# Coded Caching for Multi-level Popularity and Access

Jad Hachem, *Student Member, IEEE,* Nikhil Karamchandani, *Member, IEEE,* and Suhas Diggavi, *Fellow, IEEE*

arXiv:1404.6563v2 [cs.IT] 31 Dec 2015

**Abstract**

To address the exponentially rising demand for wireless content, use of caching is emerging as a potential solution. It has been recently established that joint design of content delivery and storage (coded caching) can significantly improve performance over conventional caching. Coded caching is well suited to emerging heterogeneous wireless architectures which consist of a dense deployment of local-coverage wireless access points (APs) with high data rates, along with sparsely-distributed, large-coverage macro-cell base stations (BS). This enables design of coded caching-and-delivery schemes that equip APs with storage, and place content in them in a way that creates coded-multicast opportunities for combining with macro-cell broadcast to satisfy users even with different demands. Such coded-caching schemes have been shown to be order-optimal with respect to the BS transmission rate, for a system with single-level content, i.e., one where all content is uniformly popular. In this work, we consider a system with non-uniform popularity content which is divided into multiple levels, based on varying degrees of popularity. The main contribution of this work is the derivation of an order-optimal scheme which judiciously shares cache memory among files with different popularities. To show order-optimality we derive new information-theoretic lower bounds, which use a sliding-window entropy inequality, effectively creating a non-cutset bound. We also extend the ideas to when users can access multiple caches along with the broadcast. Finally we consider two extreme cases of user distribution across caches for the multi-level popularity model: a single user per cache (single-user setup) versus a large number of users per cache (multi-user setup), and demonstrate a dichotomy in the order-optimal strategies for these two extreme cases.

## I. Introduction

Broadband data consumption has witnessed a tremendous growth over the past few years, due in large part to multi-media applications such as Video-on-Demand. This increased demand has been managed in the wired internet via Content Distribution Networks (CDNs), by mirroring data in various locations and in effect pushing the content closer to the end users. Wireless data consumption, driven by the increased demand for high-definition content on mobile devices, has also grown at a significant rate [1] and is testing the limits of our underlying wireless communication systems [2]. However, simply borrowing the CDN solution from wired networks and applying it to wireless systems is insufficient to solve the wireless content delivery problem. In the wired Internet, CDNs remove the bottleneck at the content distribution server by utilizing repeated demand of particular content. It has the most gains when the local communication link is not the bottleneck [3]. In wireless cellular usage, this is typically not true as the (cellular) wireless hop is a bottleneck link.

The *broadcast* nature of wireless can be used as an advantage to alleviate this problem. This, along with the emerging *heterogeneous* wireless network can be used to provide an architecture for wireless content distribution. The heterogeneous wireless network (HetNet) architecture emerging for 5G consists of a dense deployment of wireless access points (APs) with small coverage and relatively large data rates, in combination with cellular base-stations (BS) with large coverage and smaller data rates. For example, the access points could be WiFi or emerging small-cells (or femto-cells), which provide high data rate for short ranges. The consequence of this emerging architecture is that a user could potentially receive broadcast from the BS as well as connect to (several) wireless APs. Therefore, we could place caches at local APs and complement them with macro-cellular (BS) broadcast.

In this paper we study a problem based on an architecture where content is stored at multiple APs without *a priori* knowing the user requests, and the base-station broadcast is used judiciously to complement the local caching, *after* the user requests are known. This is motivated by the new approach initiated in [4], [5] where it has been shown that joint design of storage and delivery (a.k.a. "coded caching") can significantly improve content delivery rate requirements. This was enabled by content placement that creates (network-coded) multicast opportunities among users with access to different storage units, even when they have different (and *a priori* unknown) requests. This enables an examination of the optimal trade-off between cache memory size and broadcast delivery rate.

The setup studied in [4], [5] consisted of single-level content, i.e., every file in the system is uniformly demanded. However, it is well understood that content demand is non-uniform in practice, with some files being more popular than others. Motivated by this, [6], [7], [8], [9], [10], [11] considered such non-uniform content demand, following different models. In [6], [7], [8], the setup considered a single user per cache requesting a file independently and randomly according to some (arbitrary) probability distribution that represents content popularity. These works studied the trade-off between the average rate and the cache memory. A memory-sharing scheme was proposed in [6], and its achievable rate was characterized. However, from our

J. Hachem and S. Diggavi are with the Department of Electrical Engineering, University of California, Los Angeles.
N. Karamchandani is with the Department of Electrical Engineering, Indian Insitute of Technology, Bombay.
This work was supported in part by NSF grant #1423271 and a gift by Qualcomm Inc.

understanding, this scheme was not shown to be order-optimal in general.[1] In [7], a different scheme was proposed, based on a clustering of the most popular files into a single content level, which was shown to be order-optimal for Zipf-distributed content and, more recently, for arbitrary distributions in [8].

By contrast, in [9], a deterministic multi-level popularity model was introduced (simultaneous to the aforementioned other non-uniform popularity models), where content is divided into discrete levels based on popularity. In this paper, we focus on this model and study it mostly in the context where a large number of users connect to each cache ("multi-user model"), and, for each level, a fixed and *a priori* known fraction of the users per cache request files from said level. It is easy to see that, when the number of users per cache is large enough, this deterministic model will closely approximate an equivalent stochastic-demands model similar to [6], [7], [8]. We will also study the scenario where users could connect to multiple access points (caches) as well as listen to the broadcast to get the desired content. In short, the setup considered has a different popularity model and user population as well as cache access than considered in earlier literature. We will also compare the results, for this popularity model, between setups with many users per cache (multi-user setup) and a single user per cache (single-user setup).

The main contribution of this work is, for any given multi-level content popularity profile, to approximately solve the trade-off of the transmission cost at the BS with the storage cost at the APs. In addition, we also approximately solve the case where users have access to multiple APs. Finally, we study the effect of number of users per cache in the multi-level content popularity model. In particular, the following are the core technical contributions of the work:[2].

- Non-cut-set-based information-theoretic lower bounds, for the multi-user model with users accessing multiple caches, which are used to evaluate the performance of the proposed schemes.
- A memory-sharing scheme, which divides the available storage at each cache (AP) among the various popularity levels. A striking aspect of this (order-optimal, for the multi-user setup) solution is that, in some regimes, it is better to store some less popular content without completely storing the more popular content, even when cache memory is available.
- We demonstrate order-optimality of the scheme with respect to the information-theoretic lower bound that is independent of the number of popularity levels, number of users, files, and caches.
- We demonstrate that drastically different strategies are order-optimal for the multi-user and single-user setup. In the single-user case, we show that clustering the most popular levels and giving them all the memory, leaving none for the rest, is order-optimal; a strategy proposed in [7], [8]. In contrast, the multi-user case requires a complete separation of the different levels and a division of the memory between them.

The paper is organized as follows. Section II formulates the problem, describing precisely the multi-user and single-user setups. We establish some background in Section III, which enables us to state the main results in Section IV. The caching and delivery strategy for the multi-user set up, as well as corresponding lower bounds, are given in Section V, while the single-user setup is studied in Section VI. A brief discussion about the dichotomy in the two setups is given in Section VII. The paper concludes in Section VIII with a discussion and some numerical evaluations to interpret the results. Many of the detailed proofs are given in the appendices.

### A. Related work

Content caching has a rich history and has been studied extensively, see for example [12] and references therein. More recently, it has been studied in the context of Video-on-Demand systems where efficient content placement and delivery schemes have been proposed in [3], [13], [14], [15]. The impact of content popularity distributions on caching schemes has also been widely investigated, see for example [16], [17], [18].

Most of the literature has focused on wired networks and, as argued before, the solutions there do not carry directly to wireless networks. Recently, [19] proposed a caching architecture for heterogeneous wireless networks, with the small-cell or WiFi access points acting as helpers by storing part of the content. A content placement scheme is formulated and posed as a linear program. However, the (information-theoretic) optimality of such schemes was not examined in that work. Another aspect (also common to most of the papers in the content caching literature) is that the delivery phase used independent unicasts to serve the different users. The important observation to utilize broadcast to improve system performance by serving multiple users simultaneously was made in [5], [4]. They initiated the study of coded caching where joint design of storage and delivery was considered for the case with a single level of files and single cache access during delivery by proposing an order-optimal coded caching scheme. These results have been extended to online caching systems in [20], heterogeneous cache sizes [21], unequal file sizes [22], and improved converse arguments [23], [24]. Efficient coded caching schemes have been devised in [25], and the effect of finite file sizes has been investigated in [26]. Content caching and delivery has also been studied for hierarchical tree topologies [27], [28], device to device networks [29], [30], multi-server topologies [31], and heterogeneous wireless networks [10].

[1]We refer to an "order-optimal" result as one that is within a constant multiplicative factor from the information-theoretic optimum. The constant is to be independent of the number of users, caches, memory size and number of popularity levels.

[2]Shorter versions of these results have been published in [9], [10], [11].

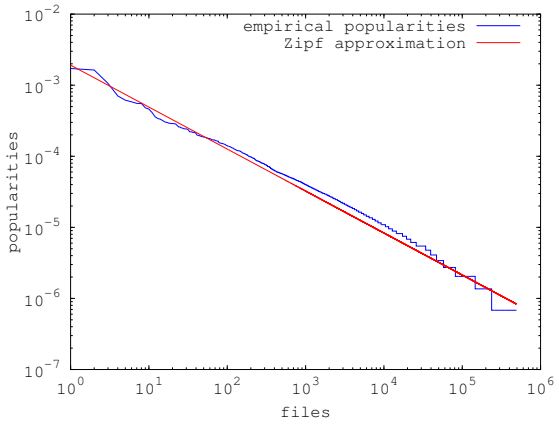


Fig. 1. Empirical popularities of some YouTube videos (based on number of views), with an approximating Zipf distribution.

Coded caching was extended to non-uniform popularity models in [6], [7], [8], [32], where the setup considered a single user per cache requesting a file independently and randomly according to some (arbitrary) probability distribution that represents content popularity. The trade-off between memory and average delivery rate was studied in these works. Our work differs from these as it uses a deterministic multi-level popularity model introduced in [9], enabling a worst-case rather than average case analysis. We analytically characterize the order-optimal splitting parameters for the memory-sharing scheme, even with user access to multiple caches. The dichotomy of order-optimal schemes between having multiple users per cache and a single user per cache is also demonstrated for this multi-level popularity model.

Other related work includes [33] which derives scaling laws for content replication in multihop wireless networks; [34] which explores distributed caching in mobile networks using device-to-device communications; [35] which studies the benefit of coded caching when the caches are distributed randomly; and [36] which explores the benefits of adaptive content placement, using knowledge of user requests.

## II. Setup, Notation, and Formulation

### A. Setup

Consider a network where a group of users request files from a server. All files are assumed to be of size $F$ bits. Prior to any user requests, a *placement phase* occurs in which information about these files is placed in the caches of $K$ access points (APs); each cache has a capacity of $MF$ bits. Then, in the *delivery phase*, users connect to the different caches, and each requests a file based on an underlying file popularity model: more popular files are more likely to be requested. The server then sends, through the base station (BS), a broadcast message of size $RF$ bits that all the users can hear. The users combine the broadcast with the contents of their cache to recover the file that they have requested. Clearly, there is a trade-off between the values of $M$ (the "cache memory") and $R$ (the "broadcast rate"): the larger the cache memory, the more information the caches can store, and hence the smaller the broadcast rate needed to serve the requests. Our goal is to characterize this trade-off.

### B. Multi-level popularity

In multi-media applications such as video-on-demand, we often find that a small number of files are requested by many more users than the rest of the files. This difference in popularity can easily influence the caching system described above. For example, when deciding what to store in the (limited-capacity) caches, one would want to give more of the cache memory to the more popular files, since they will be, on average, requested more often.

Different popularity models have been considered in the literature. The simplest model was studied in [4], [5]. In this model, all files are equally popular, meaning that there is no preference among the users to choose one file over the others. The results were of a worst-case nature, i.e., they are true for all possible (valid) combinations of user demands. While this model allowed the first approximate-optimality caching result and introduced the idea of coded caching, it is not an accurate representation of typical multi-media data. To introduce the effect of file popularity, we studied a multi-level popularity model in [9], [10], [11], in which files are divided into different popularity classes (*levels*), and files within each class are equally popular. The results are also based on a worst-case analysis. Finally, probabilistic models were also studied in the literature. In these, user demands are stochastic and follow some probability distribution, and the focus is on average results rather than worst-case ones. There has been focus on Zipf distributions [6], [7], which can be seen in examples such as YouTube videos (see Fig. 1 based on data from [37]), but also on arbitrary distributions [6], [8].

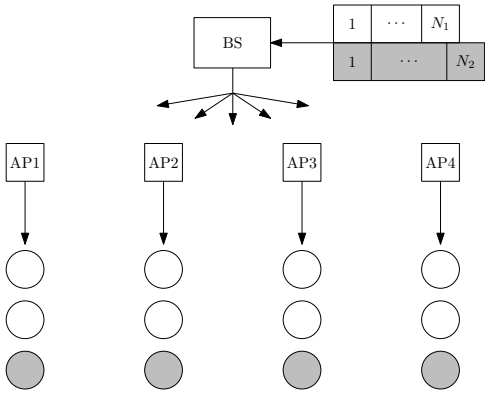


Fig. 2. Multi-user setup with $K = 4$ caches, and $L = 2$ levels with $(U_1, U_2) = (2, 1)$ users per cache. Both levels have an access degree of 1.

The popularity model that we consider in this paper is the *multi-level* model. The files are divided into $L$ *popularity levels*, such that all files in a single level are equally popular.[3] The levels consist of $N_1, \ldots, N_L$ files. Furthermore, the total number of users in the system who are requesting files from each level $i$ is *fixed* and *known* to the designer *a priori*.

To motivate the determinism in this last point, consider the following example. Suppose there are two popularity levels, and assume a stochastic-demands setup where each user is three times as likely to request a file from the first level as he is from the second. If there are $40$ users in the network, then we would expect that about $30$ of them will request a file from the first level, and $10$ from the second. By the law of large numbers, when a large number of users is present in the system, we expect a concentration of the number of users requesting files from each level around their means. Because of this concentration, the stochastic-demands model will closely resemble the determinism in the multi-level model that we adopt.

In this paper, we will use the phrase *user profile* to refer to the arrangement of users across the caches. Specifically, it is the number of users requesting a file from level $i$ at cache $k$, for every pair $(i, k)$. Moreover, we define the *user demand vector* or *user request vector* as the vector of specific files requested by each user.

### C. Multi-level access

Another aspect of the caching problem that we study is the possibility of users to access more than one cache. This introduces a third dimension to the trade-off. The more caches users can access, the more information they can gather from them for the same memory, and hence the less information they need from the broadcast transmission. However, an increased degree can be costly to users who must now establish connections with multiple APs.

We introduce this multiple-access dimension to the problem by augmenting the multi-level popularity model with a *multi-level access* aspect. Users are required to access a certain number of caches based on the popularity level of the file they request. Specifically, all users demanding a file from level $i$ must access $d_i$ caches. We call $d_i$ the *access degree* of level $i$. Because we would prefer that users connect to the nearest caches, these $d_i$ caches are consecutive, with a cyclic wrap-around for symmetry, as in Fig. 3.

### D. Number of users

As we will see later, the caching system that we have described behaves differently depending on the number of users *per cache*. In particular, it requires different strategies when this number is large or small.

We focus our attention on the two extreme cases of number of users. In one extreme, we assume that there is only one user at every cache. Recall from Section II-B that the popularity model assumes the number of users requesting a file from each level is fixed and known. However, since there is only one user per cache, we do not know beforehand which user will connect to which cache. Thus the user profile is unknown. We call this extreme the *single-user setup*. In the other extreme, the number of users per cache is large enough that the concentration of the levels requested by each user manifests itself, not only on the overall set of users, but also on the set of users connected to a single cache. In other words, we assume that the number of users requesting files from each level is fixed *per cache*. Furthermore, we assume a symmetry across the caches, so that every level is represented equally at each cache. We call this setup the *multi-user setup*.

After introducing the two setups, we now give their formal definitions.

*1) Multi-user setup:* Consider the setup shown in Fig. 3. For every level $i$ and every consecutive subset of $d_i$ caches, there are exactly $U_i$ users connecting to these caches and requesting a file from level $i$. Notice that every level is represented equally at every cache.

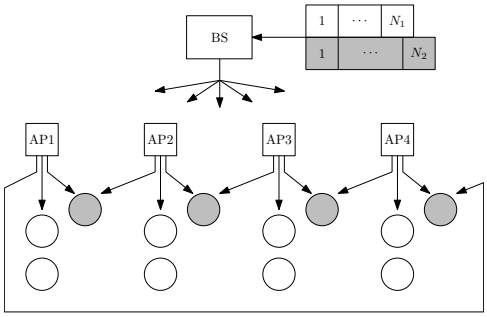


Fig. 3. Multi-user setup with $K = 4$ caches, and $L = 2$ levels with $(U_1, U_2) = (2, 1)$ users per cache, and access degrees of $(d_1, d_2) = (1, 2)$.

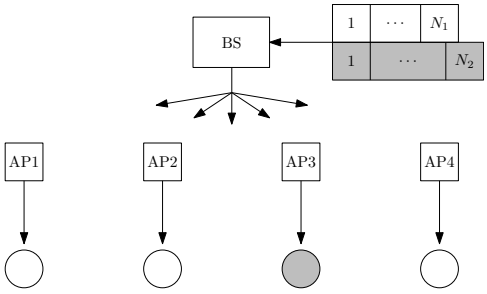


Fig. 4. Single-user setup with $K = 4$ caches, and $L = 2$ levels with $(K_1, K_2) = (3, 1)$ users.

*2) Single-user setup:* Consider now the setup in Fig. 4, depicting the other extreme. We have only one user connecting to every cache, for a total of $K$ users. The only information known *a priori* is that, for each level $i$, exactly $K_i$ out of the $K$ users will request a file from $i$. However, we do not know which users these will be.

For symmetry, we assume only single-access in the single-user setup, i.e., $d_i = 1$ for all levels $i$.

### *E. Rate-memory trade-off*

Recall that our main goal is to characterize the trade-off between the broadcast rate $R$ and the cache memory $M$. The determinism in the setups allows us to study worst-case trade-offs as opposed to average-case trade-offs.

We say that a pair $(R, M)$ is *achievable* if there exists a placement-and-delivery strategy that uses caches of memory $M$ and transmits, for *any* possible combination of user requests (valid within the popularity model), a broadcast message of rate at most $R$ that satisfies all said requests with vanishing probability as the file size $F$ grows. Our goal is to find all such achievable pairs. In particular, we wish to find the optimal rate-memory trade-off:

$$R^*(M) = \inf \{R : (R, M) \text{ is achievable}\},$$

where the minimization is done over *all* possible strategies.

### *F. Problem formulation*

The problem of finding an exact characterization of the rate-memory trade-off is difficult even for the simplest cases [4]. Therefore, in this paper, we will instead consider approximate characterizations. In particular, we wish to find an achievability strategy that results in a rate-memory trade-off $R(M)$ such that:

$$cR(M) \leq R^*(M) \leq R(M),$$

where $c$ is some constant. We say that such a strategy is *order-optimal*.

We allow $c$ to depend on only one parameter: the maximum access degree $D = \max_i d_i$. In practice, we do not expect that one user will be required to access a large number of caches, and so $D$ would be quite small. However, $c$ must be independent of all other parameters.

[3]For a discussion of the multi-level popularity model, see Section VIII-A.

TABLE I
NOTATION

| | |
|---|---|
| **All setups** | |
| $K$ | # of caches |
| $L$ | # of popularity levels |
| $N_i$ | # of files in level $i$ |
| $F$ | file size |
| $R$ | broadcast rate (normalized) |
| $M$ | cache memory (normalized) |
| $R^*(M)$ | optimal rate-memory trade-off |
| **Multi-user setup** | |
| $U_i$ | # of users per cache for level $i$ |
| $d_i$ | access degree of level $i$ |
| $D$ | maximum access degree |
| $\beta = \frac{1}{198}$ | level-separation factor |
| **Single-user setup** | |
| $K_i$ | total # of users for level $i$ |

### G. Regularity conditions

We assume the following two regularity conditions. First, for every popularity level $i$, there are more files than users. In the multi-user setup, this means:

$$\forall i, \quad N_i \geq KU_i. \tag{1}$$

In the single-user setup, we would write:

$$\forall i, \quad N_i \geq K_i. \tag{2}$$

This can be seen, for example, in video applications such as Netflix, where "files" would be video segments of a few seconds to a few minutes. To borrow an example from [27], if a database has 1000 popular movies of length 100 minutes, and each movie is divided into files (segments) of one minute each, the result is 100,000 files. It is unlikely that over 100,000 users will each be watching one of those 1000 movies at the same time.

Second, *in the multi-user setup only*, we assume that no two levels have very similar popularities. The popularity of a level can be written as the number of users per file of the level, i.e., as $U_i/N_i$. Hence, if $i$ is a more popular level than $j$, the regularity condition states:[4]

$$\sqrt{\frac{U_i/N_i}{U_j/N_j}} \geq \frac{D}{\beta} = 198D, \tag{3}$$

where $\beta = \frac{1}{198}$ is called the *level-separation factor*. The reasoning behind this condition is that, if it did not hold for some levels $i$ and $j$, then we can think of them as essentially one level with $N_i + N_j$ files and $U_i + U_j$ users per cache. The resulting popularity $\frac{U_i+U_j}{N_i+N_j}$ would be close to both $U_i/N_i$ and $U_j/N_j$.

### H. Notation table

For reference, we present in TABLE I all the notation that we use in this paper.

## III. PRELIMINARIES

Traditional caching only uses multiple-unicast transmissions from the server to the users. As a result, the total transmission size was proportional to the number of users, for any value of the cache memory. Coded caching, initially introduced in [4], brought a drastic improvement by eliminating the dependence of the transmission size on the number of users (except for very small cache memory). This technique was shown to be approximately optimal in [4] (a centralized version) and in [5] (a decentralized version), under a setup with a single level of popularity and a single user at every cache, with a single-access structure. We will refer to this setup as the Basic Setup, because it will form the basis of our main analysis.

The scheme for the Basic Setup, in its decentralized form (on which will we henceforth focus) consists in placing a random sampling of bits from all files in every cache, independently. Consequently, there will be some overlap, but also some differences, in the bits present in every cache. The BS then transmits linear combinations of these bits, taking advantage of the overlaps as well as the differences, so that the same linear combination can be useful for multiple users at once. The resulting rate-memory trade-off is given in the following Lemma.

[4]As we will see in later sections, the variables $U_i$ and $N_i$ will often appear inside square roots. For this reason, phrasing the regularity condition using the square roots is more useful.

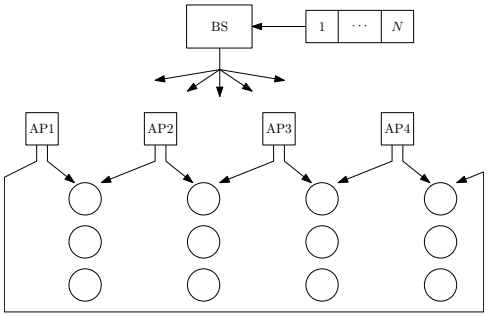


Fig. 5. Generalized single-level setup with $U = 3$ users per cache, and an access degree $d = 2$.

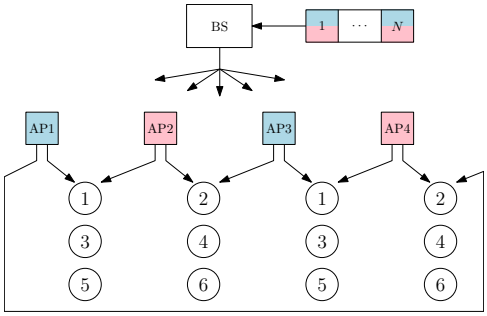


Fig. 6. An illustration of the scheme used on the example from Fig. 5. Caches are colored into $d = 2$ colors, and the files are divided and colored with the same colors. In parallel, the users are divided into $dU = 6$ groups, where users fro m the same group have no overlapping caches.

**Lemma 1** (Rate for the Basic Setup [4])**.** *For a single-level caching system with $K$ caches, $N$ files, a single user per cache with an access degree of 1, and a cache memory of $M$, the following rate is achievable:*

$$R^0(M, K, N) = \min\left\{\frac{N}{M}, K\right\} \cdot \left(1 - \frac{M}{N}\right).$$

*Furthermore, this rate is within a constant of the optimum.*

Notice that, when $M > N/K$, then the rate becomes $(N/M - 1)$, removing all dependence on the number of users. Under traditional caching, the rate would have been $K(1 - M/N)$.

### *A. Generalizing to multi-user, multi-access*

In the multi-level setups, every level can have more than one user per cache, and can also have an access degree that is larger than 1. Thus, the first step towards solving the multi-level problem is to generalize the single-level problem to one where these two parameters are no longer 1. Fig. 5 illustrates such a setup.

We adopt the following strategy. We first color all the caches into $d$ colors, such that every user is connected to exactly one cache of every color. In parallel, we split every file into $d$ equal subfiles, and color each subfile using the same colors as the caches. Next, we group all the $KU$ users into $dU$ groups of $K/d$ users each, such that no two users in the same group share any cache. The end result is illustrated in Fig. 6.

In the placement phase, we consider each of the $d$ colors separately. The caches of a certain color perform a simple placement on the subfiles of the same color. Since each file is split into $d$ equal subfiles, every cache can hold $dM$ subfiles, the equivalent of $M$ files.

In the delivery phase, each of the $dU$ groups is considered separately. For every group and every color, we have a subsystem where $K/d$ users are each requesting a subfile of size $F/d$ bits from one like-colored cache of size $dM$ subfiles. The total broadcast size is thus:

$$RF = dU \cdot d \cdot R^0(dM, K/d, N) \cdot (F/d) \text{ bits.}$$

Note that the memory given to the $R^0$ function is $dM$. This is because the "files" in each subsystem have a size of $F/d$ bits, and thus the total cache size of $MF$ bits must be normalized by $F/d$ bits.

Using the above equation with Lemma 1, we get the following theorem.

**Theorem 1** (Single-level rate-memory trade-off)**.** *Given a single-level caching system, with $N$ files, $K$ caches, $U$ users at each cache with access degree $d$, and a cache memory of $M$, the following rate is achievable:*

$$R^{SL}(M, K, N, U, d) = U \cdot \min\left\{\frac{N}{M}, K\right\} \cdot \left(1 - \frac{dM}{N}\right). \tag{4}$$

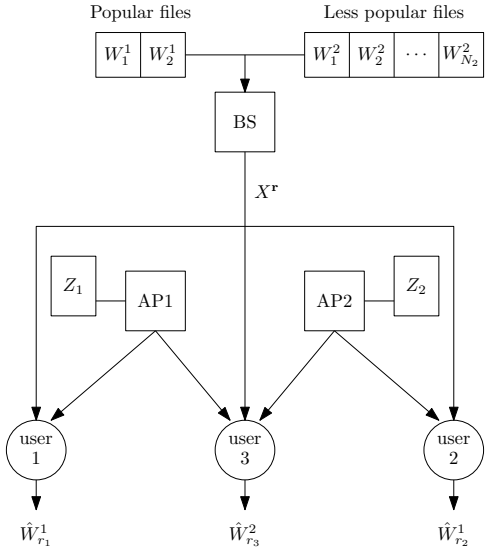


Fig. 7. Small example that illustrates multi-level popularity and access.

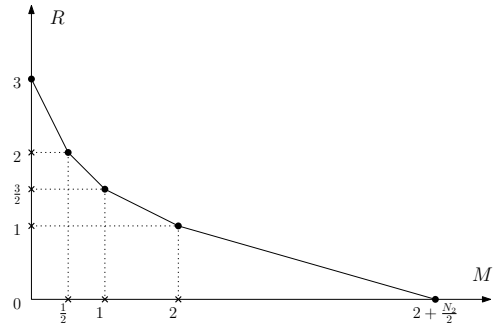


Fig. 8. Optimal rate-memory trade-off $R^*(M)$ for the small example.

While, in the single-access case, a cache memory $M < N$ meant a non-zero transmission rate, adding the multi-access aspect results in a zero achievable rate for the smaller memory value of $N/d$. Intuitively, at $M = N/d$, it is as though we apply an erasure-correcting code on all the files and spread it across the caches, such that any $d$ caches can reproduce all files.

The results presented in this section will be key to our solution of the multi-level caching problem. Indeed, in all that will follow, we use the above-described (decentralized) coded caching scheme as a black box that gives a rate $R^{\text{SL}}(M, K, N, U, d)$ for input parameters $M$, $K$, $N$, $U$, and $d$.

### *B. A small multi-level example with exact characterization*

In order to illustrate the general multi-level problem, we will here present a small example that combines both multi-level popularity and multi-level access. We give, for this example, an exact characterization of the rate-memory trade-off.

Consider the setup in Fig. 7. The server holds files from $L = 2$ popularity levels. The first level has $N_1 = 2$ files, and the second level has $N_2 \geq 4$ files. There are two APs, each equipped with a cache of memory (i.e., size normalized by file size) $M$. There is one user accessing each cache and requesting a file from level 1 (users 1 and 2), and a third user accessing both caches and requesting a file from level 2 (user 3).

**Theorem 2** (Exact characterization for the small example)**.** *For the setup shown in Fig. 7, the optimal rate-memory trade-off is plotted in Fig. 8 and is characterized by:*

$$R^*(M) = \max\left\{3 - 2M, \frac{5}{2} - M, 2 - \frac{1}{2}M, 1 - \frac{M-2}{N_2/2}\right\}.$$

The proof of Theorem 2 is given in Appendix C. However, for illustration and to gain intuition about the general problem, we will here briefly discuss the achievability scheme for two values of the cache memory $M$: $M = 1$ (for which $R = 3/2$ is achievable) and $M = 1/2$ (for which $R = 2$ is achievable).

Suppose that $M = 1$, so that each cache can hold the equivalent of one file. We first split each file in level 1 into two equal parts: $W_n^1 = (W_{n,a}^1, W_{n,b}^1)$, for $n = 1, 2$. Now, each cache exclusively stores one half of each popular file, and stores nothing

from level 2. Thus the first cache will contain $(W^1_{1,a}, W^1_{2,a})$ and the second cache will contain $(W^1_{1,b}, W^1_{2,b})$. When the users make their requests, the BS transmits $W^2_{r_3}$ completely for user 3, and sends a coded transmission $(W^1_{r_1,b} \oplus W^1_{r_2,a})$ for users 1 and 2 together. Combining the transmission with the contents of their respective caches, each of users 1 and 2 can recover the file that they have requested. In total, the BS would have transmitted one complete file ($W^2_{r_3}$), plus the equivalent of one half-file (the linear combination), for a total rate of $R = 3/2$.

Suppose now that $M = 1/2$, i.e., each cache can only hold the equivalent of half a file. We again split the two level-1 files just like in the previous case. However, this time the first cache stores $(W^1_{1,a} \oplus W^1_{2,a})$ while the second cache stores $(W^1_{1,b} \oplus W^1_{2,b})$. When the users make their requests, the BS again transmits $W^2_{r_3}$ to serve user 3, but also sends $W^1_{r_1,b}$ and $W^1_{r_2,a}$ for users 1 and 2. This allows them to recover the file that they have requested by combining the transmission with the side-information available at their caches. Since the BS has transmitted one complete file and two half-files, the total rate is $R = 2$.

While, in this small example, an exact characterization of the rate-memory trade-off was found, this is difficult in general. For this reason, we focus our attention on order-optimality results as stated in Section II-F.

## IV. Main Results

In this section, we provide the approximately optimal rate-memory trade-off for each of the two setups (multi-user and single-user). We discuss how each such trade-off is achieved.

### A. Multi-user setup

The placement-and-delivery strategy that we adopt for the multi-user setup is a *memory-sharing* strategy. It consists of dividing the cache memory between all the $L$ levels, and then treating each level as a separate caching sub-system, with the reduced memory. In other words, we give level $i$ a memory $\alpha_i M$, where $\alpha_i \in [0,1]$ and $\sum_i \alpha_i = 1$, and we then apply a single-level placement-and-delivery strategy for this level on this $\alpha_i M$ memory, separately from the other levels. The total rate for this scheme is thus:

$$R^{\text{MU}}\left(M, K, \{N_i, U_i, d_i\}_i\right) = \sum_{i=1}^{L} R^{\text{SL}}(\alpha_i M, K, N_i, U_i, d_i), \tag{5}$$

where $R^{\text{SL}}(\cdot)$ is defined in (4).

By optimizing the overall rate over the memory-sharing parameters $\{\alpha_i\}_i$, we establish a memory allocation which we will show is order-optimal. At a high level, this allocation is done by partitioning the popularity levels into three sets: $H$, $I$, and $J$. The levels in $H$ have such a small popularity that they will get no cache memory. On the opposite end of the spectrum, the most popular levels are assigned to $J$ and are given enough cache memory to completely store all their files in every cache. Finally, the rest of the levels, in the set $I$, will share the remaining memory among themselves, obtaining some non-zero amount of memory but not enough to store all of their files.

Our choice of the $(H, I, J)$ partition and corresponding memory assignments are discussed in Section V-A. This choice results in the following achievable rate.

**Theorem 3.** *Given a multi-user caching setup, with $K$ caches, $L$ levels, and, for each level $i$, $N_i$ files and $U_i$ users per cache with access degree $d_i$, and a cache memory of $M$, the following rate[5] is achievable:*

$$R^{MU}(M) \approx \sum_{h \in H} KU_h + \frac{\left(\sum_{i \in I} \sqrt{N_i U_i}\right)^2}{M - \sum_{j \in J} N_j/d_j} - \sum_{i \in I} d_i U_i,$$

*where $(H, I, J)$ is a particular type of partition of the set of levels called an $M$-feasible partition.*[6]

Intuitively, since a level $h \in H$ receives no cache memory, all requests from its $KU_h$ users must be handled directly from the broadcast. Since, by regularity condition (1), we have $N_i \geq KU_i$ for all levels $i$, then in the worst case a total of $KU_h$ distinct files must be completely transmitted for the users requesting files from level $h$. The users in set $J$ require no transmission as the files are completely stored in all the caches; however, it does affect the rate through the memory available for levels in $I$. This is apparent in the expression $M - \sum_{j \in J} N_j/d_j$. Finally, the levels in $I$, having received some memory, require a rate that is inversely proportional to the effective memory and that depends on the level-specific parameters $N_i$, $U_i$, and $d_i$.

The structure of the $(H, I, J)$ partition that we have chosen allows us to efficiently compute it for all values of the cache memory $M$. Indeed, we provide an algorithm in Section V-A that can find this partition, as well as the corresponding memory-sharing parameters $\alpha_i$, for every $M$ in $\Theta(L^2)$ running time. Briefly, as $M$ is increased, levels get "promoted" from the set $H$ to $I$ to $J$. The sequence of these promotions is directly determined by the popularity of the levels.

[5]This expression of the rate is a slight approximation that we use here for simplicity as it is more intuitive. An exact and complete description of the achievable rate is given in Section V-A.

[6]This type of partition is defined in Definition 1 and elaborated upon in Section V-A.

We now discuss the order-optimality of the memory-sharing scheme in the multi-user setup. We develop new, non-cut-set lower bounds on the optimal rate, which use sliding-window entropy inequalities [38], and show that the scheme achieves a rate that is within a constant factor of the optimal. Note that this constant is independent of all the problem parameters except the largest AP access degree $D$.

**Theorem 4.** *For all valid values of the problem parameters $K$, $L$, $\{N_i, U_i, d_i\}_i$, and $M$, we have:*

$$\frac{R^{MU}(M)}{R^*(M)} \leq cD,$$

*where $R^{MU}(M)$ is the rate achieved by memory-sharing, $R^*(M)$ is the optimal rate over all strategies, $D$ is the largest AP access degree $D = \max_i d_i$, and $c = 9909$ is a constant (independent of all problem parameters).*

The gap between the rate achieved by the memory-sharing strategy and the optimal rate is linear in $D$. As we have argued earlier, we would not expect a situation where one user connects to too large a number of APs, and so $D$ can be thought of as a constant.

The lower bounds needed to prove Theorem 4 have to include the effect of all the popularity levels on the transmission rate. However, these effects can be very different, especially when some files are much more popular than others, something that cut-set bounds alone cannot account for. Using sliding-window subset entropy inequalities [38], we can combine multiple cut-set bounds that correspond to the different levels, without making any assumptions on the achievability scheme. The resulting bounds bring out the necessity for memory-sharing.

Furthermore, the constant $c$ in Theorem 4 is rather loose so as to simplify the analysis. Numerics show that, in practice, this constant is much smaller. For example, if we have $K = 20$ caches, $L = 3$ popularity levels consisting of $(N_1, N_2, N_3) = (200, 20\,000, 800\,000)$ files, $(U_1, U_2, U_3) = (10, 5, 1)$ users per cache, and access degrees of $(d_1, d_2, d_3) = (1, 1, 1)$, then the gap is less than $6.8$.

### *B. Single-user setup*

In the single-user setup, the scheme that we propose is quite different. Instead of separating the levels, we *cluster* a subset of them into a super-level that will be treated as essentially one level. Specifically, we partition the levels into two subsets: $H'$ and $I'$. The set $I'$ will be clustered into one super-level, and all of the memory $M$ will be given to it, while $H'$ will be given no memory.

To understand how to choose $H'$ and $I'$, consider the following rough analysis. Suppose that all levels except one (let's call it $j$) have been split into $H'$ and $I'$. Then, the rate, using Theorem 1, would be:

$$\begin{aligned} R &= R^{\text{SL}}(0, \textstyle\sum_{h\in H'} K_h, \sum_{h\in H'} N_h, 1, 1) \\ &\quad + R^{\text{SL}}(M, \textstyle\sum_{i\in I'} K_i, \sum_{i\in I'} N_i, 1, 1) \\ &\approx \sum_{h\in H'} K_h + \frac{\sum_{i\in I'} N_i}{M}. \end{aligned}$$

If we were to add level $j$ to $H'$, that would result in the addition of a $K_j$ term, since all $K_j$ requests would be completely served by the broadcast. On the other hand, if it is added to $I'$, then we would get an additional $N_j/M$ term, since the total number of files in $I'$ would increase by $N_j$. Clearly, it is beneficial to choose the smaller of the two quantities.

Though the above analysis is rough, its main idea still holds. In general, we choose the partition $(H', I')$ as follows:

$$H' = \left\{ h \in \{1, \ldots, L\} : M < \frac{N_h}{K_h} \right\}; \quad I' = (H')^c. \tag{6}$$

Then, by giving all of the memory to $I'$, we can apply a single-level caching-and-delivery scheme to obtain the rate in the following theorem.

**Theorem 5.** *Consider the multi-level, single-user setup with $L$ levels, $N_i$ files and $K_i$ users for each level $i$, and cache memory $M$. Then, the following rate is achievable:*

$$R^{SU}(M) = \sum_{h\in H'} K_h + \max\left\{ \frac{\sum_{i\in I'} N_i}{M} - 1\,,\, 0 \right\},$$

*where $H'$ and $I'$ are as defined in* (6).

This scheme turns out to be order-optimal, as we state in the next theorem.

**Theorem 6.** *The rate achieved by the clustering strategy in Theorem 5 is within a constant multiplicative factor of the information-theoretic optimum:*

$$\frac{R^{SU}(M)}{R^*(M)} \leq 72,$$

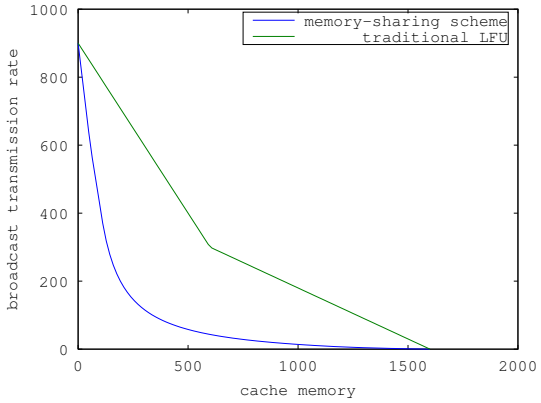


Fig. 9. Comparison of the memory-sharing scheme with traditional LFU, in the multi-user setup.

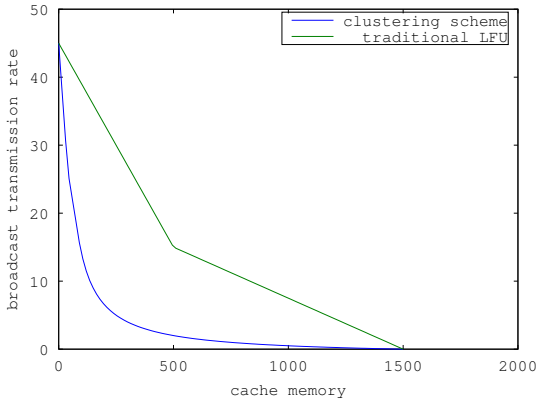


Fig. 10. Comparison of the clustering scheme with traditional LFU, in the single-user setup.

*where $R^{SU}(M)$ is the rate achieved by clustering, and $R^*(M)$ is the information-theoretically optimal rate.*

Unlike in the multi-user case, cut-set bounds are sufficient to show order-optimality in this case. Indeed, a single cut-set bound allows us to capture the fact that the user profile (i.e., the level of the file requested at each cache, as defined in Section II-B) is not determined beforehand. At the same time, it brings out the necessity of clustering levels by mixing their demands. As before, however, these bounds do not make any assumptions on the achievability strategy.

The scheme suggested in Theorem 5 is similar to the results in [7] for Zipf popularity distributions and in [8] for arbitrary distributions. However, this is done for the multi-level popularity model, and Theorem 6 establishes a universal approximation for worst case rate, rather than average rate.

### C. Comparison with LFU

One of the most common caching schemes conventionally in use is the Least-Frequently-Used (LFU) scheme. This strategy stores the most popular files only, as many as the caches can hold. In traditional LFU, no coding is done in the delivery phase. We give an example for each of the two setups, to show how the respectively chosen schemes are superior to LFU in each context.

*1) Multi-user example:* Consider two levels ($L = 2$), with $K = 30$, $(N_1, N_2) = (600, 1000)$, $(U_1, U_2) = (20, 10)$, and $(d_1, d_2) = (1, 1)$. The rates achieved by the memory-sharing strategy and traditional LFU are plotted against memory $M$ in Fig. 9. Memory-sharing performs up to 29 times better than LFU in this example.

*2) Single-user example:* Consider two levels ($L = 2$), with $(N_1, N_2) = (500, 1000)$ and $(K_1, K_2) = (30, 15)$. The rates achieved by the clustering strategy and traditional LFU are plotted against memory $M$ in Fig. 10. Clustering performs up to 22 times better than LFU in this example.

## V. The Multi-User Setup

### A. Caching-and-delivery strategy: memory-sharing

In this section, we describe the strategy used to achieve the rate approximated by the expression in Theorem 3. Moreover, we will give an exact upper bound on the achieved rate.

We will proceed in two steps. First, we will discuss the $(H, I, J)$ partition of the set of levels as first described in Section IV-A. This will be accompanied by an explanation of the memory-sharing parameters $\alpha_i$, which indicate what fraction of memory each level gets. Second, we analyze the *individual rate* achieved by every level $i$, after allocating $\alpha_i M$ memory to it, and combine all levels to produce the total rate achieved by the scheme.

While we actually define the $(H, I, J)$ partition based on the problem parameters, *and then* choose the $\alpha_i$ values accordingly, in this paragraph we will proceed in the opposite order so that we expose the intuition behind the choices. As discussed in Section IV-A, the strategy involves finding a good partition $(H, I, J)$ of the set of levels, such that levels in $H$ are given no memory, levels in $J$ are given maximal memory, and levels in $I$ share the rest. Thus, for all levels $h \in H$, we will assign $\alpha_h M = 0$. Similarly, every level $j \in J$ will receive $\alpha_j M = N_j/d_j$, since that is the amount of memory needed to completely store level $j$ and hence to require no BC transmission (i.e., $R^{\text{SL}} = 0$; see Theorem 1). What is left is to share the remaining memory $\left(M - \sum_{j\in J} N_j/d_j\right)$ among the levels in $I$. More popular files should get more memory, and the popularity of a level $i$ is proportional to $U_i/N_i$. Thus, we choose to give level $i$ a memory roughly $\alpha_i M \propto N_i \cdot \sqrt{U_i/N_i}$ (hence the memory *per file* is proportional to $\sqrt{U_i/N_i}$).[7]

Intuitively, we want to choose $H$, $I$, and $J$ such that the above values of $\alpha_i$ are valid, i.e., $\alpha_i \in [0, 1]$ for all $i$. Based on the intuition, we get the partition described below.

**Definition 1** ($M$-feasible partition)**.** *For any cache memory $M$, an $M$-feasible partition $(H, I, J)$ of the set of levels is a partition that satisfies:*

$$\begin{aligned}
&\forall h \in H, && \tilde{M} < \frac{1}{K}\sqrt{\frac{N_h}{U_h}};\\
&\forall i \in I, && \frac{1}{K}\sqrt{\frac{N_i}{U_i}} \le \tilde{M} \le \left(\frac{1}{d_i} + \frac{1}{K}\right)\sqrt{\frac{N_i}{U_i}};\\
&\forall j \in J, && \left(\frac{1}{d_j} + \frac{1}{K}\right)\sqrt{\frac{N_j}{U_j}} < \tilde{M},
\end{aligned}$$

*where $\tilde{M} = (M - T_J + V_I)/S_I$, and, for any subset $A$ of the levels:*

$$S_A = \sum_{i\in A}\sqrt{N_i U_i}; \quad T_A = \sum_{i\in A}\frac{N_i}{d_i}; \quad V_A = \sum_{i\in A}\frac{N_i}{K}.$$

We stress again that, while we used our own $\alpha_i$ values to determine $(H, I, J)$, this was done only for the intuition behind the choice. The partition itself is defined solely based on the problem parameters, and not on our strategy.

Notice in Definition 1 how the different sets are largely determined by the quantity $\sqrt{N_i/U_i}$, for each level $i$. This matches the idea that the most popular levels (i.e., those with the *smallest* $N_i/U_i$) will be in $J$, while the least popular levels (those with the *largest* $N_i/U_i$) will go to the set $H$.

After choosing an $M$-feasible partition, we share the memory among the levels using the following (precise) values of $\alpha_i$:

$$\begin{aligned}
&\forall h \in H, && \alpha_h M = 0;\\
&\forall i \in I, && \alpha_i M = \sqrt{N_i U_i}\cdot\tilde{M} - N_i/K;\\
&\forall j \in J, && \alpha_j M = N_j/d_j.
\end{aligned} \tag{7}$$

For completeness, the following proposition states the validity of this choice of memory-sharing parameters.

**Proposition 1.** *The values of the memory-sharing parameters defined in* (7) *satisfy:*

1) $\alpha_i \ge 0$ *for all $i$;*
2) $\sum_i \alpha_i = 1$*;*
3) $\alpha_i M \le N_i/d_i$ *for all $i$.*

*Note that points 1 and 2 imply $\alpha_i \in [0, 1]$.*

*Proof:* All three results follow directly from Definition 1. ∎

The structure of the solution described above allows us to efficiently compute the $\alpha_i$ values. Indeed, Algorithm 1 finds $\alpha_i$ for all levels $i$ and *for all memory values* $M$ in $\Theta(L^2)$ running time, where $L$ is the total number of levels. While a detailed description and analysis of the algorithm is given in Appendix A-C, we briefly go over it in this section. It proceeds in three main steps. In the first step, the algorithm identifies the sequence of $(H, I, J)$ partitions that will occur as $M$ increases, using only the problem parameters. For example, if there are two levels, there are two possible sequences, denoted below as S1 and S2:

[7]The square root comes from minimizing an inverse function of $\alpha_i$.

| | | $M=0$ | | | | large $M$ |
|---|---|---|---|---|---|---|
| S1 | $H$ | $\{1,2\}$ | $\{2\}$ | $\emptyset$ | $\emptyset$ | $\emptyset$ |
| | $I$ | $\emptyset$ | $\{1\}$ | $\{1,2\}$ | $\{2\}$ | $\emptyset$ |
| | $J$ | $\emptyset$ | $\emptyset$ | $\emptyset$ | $\{1\}$ | $\{1,2\}$ |
| S2 | $H$ | $\{1,2\}$ | $\{2\}$ | $\{2\}$ | $\emptyset$ | $\emptyset$ |
| | $I$ | $\emptyset$ | $\{1\}$ | $\emptyset$ | $\{2\}$ | $\emptyset$ |
| | $J$ | $\emptyset$ | $\emptyset$ | $\{1\}$ | $\{1\}$ | $\{1,2\}$ |

Notice that the $M$-feasible partition is the same throughout an entire interval of values for $M$. In fact, there are only $2L$ non-trivial intervals: each interval is distinguished from the previous one by the "promotion" of a level from $H$ to $I$ or from $I$ to $J$. As a result, computing only $2L$ intervals (actually, $(2L+2)$ intervals, to include the boundary cases) is enough to determine the $M$-feasible partition for *all* values of $M \geq 0$.

While Step 1 determines which of the above two sequences (S1 and S2) will occur, it does not determine the boundaries of the different regimes, i.e., the memory values at which the $(H, I, J)$ partition changes (except for trivial boundaries like $M = 0$). The second step computes these boundaries, using not only the problem parameters, but also the partitions themselves.

For every $M$, the corresponding $M$-feasible partition is determined by the algorithm based on the boundaries calculated in Step 2. For example, suppose the algorithm decided that sequence S1 should be in use, and that the boundaries of $(H, I, J) = (\emptyset, \{1, 2\}, \emptyset)$ are some values $m_1$ and $m_2$. Then, any $M \in [m_1, m_2]$ has $(\emptyset, \{1, 2\}, \emptyset)$ as its $M$-feasible partition.

**Algorithm 1** An algorithm that constructs an $M$-feasible partition for all $M$.

**Require:** Number of caches $K$ and parameters $\{N_i, U_i, d_i\}_i$ for $i = 1, \ldots, L$.
**Ensure:** An $M$-feasible partition for all $M$.

**for all** $i \in \{1, \ldots, L\}$ **do**
  $\tilde{m}_i \leftarrow (1/K)\sqrt{N_i/U_i}$
  $\tilde{M}_i \leftarrow (1/d_i + 1/K)\sqrt{N_i/U_i}$
**end for**
$(x_1, \ldots, x_{2L}) \leftarrow \mathrm{sort}(\tilde{m}_1, \ldots, \tilde{m}_L, \tilde{M}_1, \ldots, \tilde{M}_L)$.

*Step 1: Determine $(H, I, J)$ for each interval $(x_t, x_{t+1})$.*
Set $H_0 \leftarrow \{1, \ldots, L\}$, $I_0 \leftarrow \emptyset$, $J_0 \leftarrow \emptyset$.
**for** $t \in \{1, \ldots, 2L\}$ **do**
  **if** $x_t = \tilde{m}_i$ for some $i$ **then**
    *Promote level $i$ from $H$ to $I$*
    $H_t \leftarrow H_{t-1} \setminus \{i\}$
    $I_t \leftarrow I_{t-1} \cup \{i\}$
    $J_t \leftarrow J_{t-1}$
  **else if** $x_t = \tilde{M}_i$ for some $i$ **then**
    *Promote level $i$ from $I$ to $J$*
    $H_t \leftarrow H_{t-1}$
    $I_t \leftarrow I_{t-1} \setminus \{i\}$
    $J_t \leftarrow J_{t-1} \cup \{i\}$
  **end if**
**end for**

*Step 2: Compute the limits of the intervals as $[Y_t, Y_{t+1})$.*
**for all** $t \in \{1, \ldots, 2L\}$ **do**
  $Y_t \leftarrow x_t \cdot S_{I_t} + T_{J_t} - V_{I_t}$
**end for**
$Y_{2L+1} \leftarrow \infty$ ▷ For convenience

*Step 3: Determine the $M$-feasible partition for all $M$.*
**for all** $t \in \{1, \ldots, 2L\}$ **do**
  Set $(H_t, I_t, J_t)$ as the $M$-feasible partition of *all* $M \in [Y_t, Y_{t+1})$
**end for**

The following lemma presents two important properties of $M$-feasible partitions. It will be proved in Appendix A-C.

**Lemma 2.** *For any cache memory $M$, an $M$-feasible partition $(H, I, J)$ always exists. Furthermore, the set $I$ is never empty as long as $M \leq \sum_i N_i/d_i$, i.e., as long as individual caches do not have enough memory to store everything.*

To properly analyze the achievable rate, we need to look more closely at the set $I$. In the single-level, single-access scenario in [4], [5], three regimes were identified, and they were analyzed separately. Generalizing to the single-level, multi-access case, these regimes are: when $M < N/K$, when $M > cN/d$ for some constant $c \in (0,1)$, and the intermediate case. We identify three similar regimes for each level in $i$. Formally, define:

$$\begin{aligned} I_0 &= \left\{ i \in I : \tilde{M} < \frac{2}{K}\sqrt{\frac{N_i}{U_i}} \right\}; \\ I_1 &= \left\{ i \in I : \tilde{M} > \left(\frac{\beta}{d_i} + \frac{1}{K}\right)\sqrt{\frac{N_i}{U_i}} \right\}; \\ I' &= I \setminus (I_0 \cup I_1), \end{aligned} \tag{8}$$

By choosing the $\alpha_i$ values in (7), these definitions are equivalent to: $I_0$ is the set of levels $i$ such that $\alpha_i M < N_i/K$; $I_1$ is such that $\alpha_i M > \beta N_i/d_i$ for all $i \in I_1$, where $\beta$ is the level-separation factor as defined in Regularity Condition (3); and $I'$ consists of the remaining levels.

When $K \geq D/\beta$, then $I_0$, $I'$, and $I_1$ are mutually exclusive and form a partition of $I$. For convenience, we call the resulting partition $(H, I_0, I', I_1, J)$ a *refined $M$-feasible partition*.[8]

What follows is an important statement regarding the set $I_1$.

**Proposition 2** (Size of $I_1$)**.** *In any refined $M$-feasible partition $(H, I_0, I', I_1, J)$ as defined in Definition 1 and* (8)*, the set $I_1$ contains at most one element.*

*Proof:* Suppose that there exist two level $i, j \in I_1$ (and possibly others). We will show that this violates regularity condition (3).

Suppose without loss of generality that $i$ is more popular than $j$. Since $i, j \in I_1$, then, by Definition 1 and (8):

$$\left(\frac{\beta}{d_j} + \frac{1}{K}\right)\sqrt{\frac{N_j}{U_j}} < M \leq \left(\frac{1}{d_i} + \frac{1}{K}\right)\sqrt{\frac{N_i}{U_i}}.$$

However, this means:

$$\sqrt{\frac{U_i/N_i}{U_j/N_j}} < \frac{1/d_i + 1/K}{\beta/d_j + 1/K} < \frac{d_j}{\beta d_i} \leq \frac{D}{\beta},$$

which contradicts regularity condition (3). ∎

Using the definition of a refined $M$-feasible partition, and the values of $\{\alpha_i\}_i$, we give upper bounds on the rates achieved individually for each level.

**Lemma 3.** *Given a refined $M$-feasible partition $(H, I_0, I', I_1, J)$, the individual rates of the levels are upper-bounded by:*

$$\begin{aligned} &\forall h \in H, && R_h(M) = KU_h; \\ &\forall i \in I_0 \cup I', && R_i(M) \leq \frac{2S_I\sqrt{N_i U_i}}{M - T_J + V_I}; \\ &\forall i_1 \in I_1, && R_{i_1}(M) \leq \frac{1}{\beta} d_{i_1} U_{i_1} \left(1 - \frac{M - T_J}{N_{i_1}/d_{i_1}}\right) \\ & && \qquad + \frac{1}{\beta} d_{i_1} U_{i_1} \frac{S_{I_0} + S_{I'}}{\sqrt{N_{i_1} U_{i_1}}}; \\ &\forall j \in J, && R_j(M) = 0. \end{aligned}$$

*The total achieved rate is then: $R(M) = \sum_{i=1}^{L} R_i(M)$.*

We relegate the proof of Lemma 3 to Appendix A-C. What follows is a brief explanation of the individual rates.

- Since levels $h \in H$ get no memory, no information about their files can be stored in the caches. Hence, the server will have to transmit a complete copy of every file requested by the $KU_h$ users in level $h$. Therefore, $R_h(M) = KU_h$ in the worst case.
- Since levels $j \in J$ get maximal memory, any file requested from $j$ can be fully recovered using the caches. Thus there is no need for the server to send anything for $j$, and hence $R_j(M) = 0$.
- Finally, the levels $i \in I$ get enough of the remaining memory $M - T_J$ so that they behave as in Theorem 1. However, since $\alpha_i M \approx \sqrt{N_i U_i}/S_I \cdot (M - T_J)$, we get a rate of:

$$R_i(M) \approx \frac{N_i U_i}{\alpha_i M} - d_i U_i \approx \frac{S_I\sqrt{N_i U_i}}{M - T_J} - d_i U_i.$$

[8]The fact that this is a partition is important for the case $K \geq D/\beta$. In the opposite case, this does not necessarily hold but it will not matter anyway.

Notice how $R_i(M)$ is inversely proportional to the memory remaining after storing $J$. This behavior is captured in the expressions for $I_0 \cup I'$. However, for the levels in $i_1 \in I_1$, which get almost, but not quite, maximal memory, the individual rate $R_{i_1}(M)$ behaves more closely to a linear function of the memory.

### *B. Outer bounds*

In the single-level setup introduced in [4], cut-set lower bounds on the optimal rate were sufficient for proving order-optimality of the scheme. The idea was to choose cuts that include a certain number of caches $s$, as well as enough broadcast messages to decode almost all files when combined with the $s$ caches. The choice of $s$ depended most importantly on the size of the memory $M$. For example, for very large $M$, the bounds would reflect the fact that a single cache should contain enough information so that it need only be combined with a small broadcast message to decode one file. In contrast, for small $M$, the bounds would instead show that a large broadcast message is needed to serve all $s$ users.

In our multi-level situation, the goal of our lower bounds is to show that different popularity levels require different amounts of memory, ranging from zero to maximal memory. However, a single cut-set bound would force all levels to abide by the same number of caches $s$, a choice that would not reflect the plurality of memory allocation. New lower bounds were thus required.

The lower bounds that we use allow choosing for each level $i$ the number of caches $s_i$ that corresponds to the amount of memory we expect it to get. In essence, each level $i$ will exclusively be part of a cut-set bound with $s_i$ caches, which will give a lower bound that loosely corresponds to the individual rate $R_i$ for this level.

The challenge is to incorporate all of these cut-set bounds together, such that the same caches appear in cut-set bounds for all levels. This process is described next. We start with a cut-set bound for the level with the smallest $s_i$. Using Fano's inequality, we can derive lower bounds on its individual rate. We then proceed to lower-bound the current quantity using a cut-set bound for the level with the next smallest $s_i$, and repeat the process. The transition from one cut-set bound to the next can be achieved using sliding-window subset entropy inequalities [38, Theorem 3]:[9]

**Lemma 4** (Sliding-window subset entropy inequality [38, Theorem 3])**.** *Given $K$ random variables $(Y_1, \ldots, Y_K)$, we have, for every $s \in \{1, \ldots, K-1\}$:*

$$\frac{1}{s}\sum_{i=1}^{K} H\left(Y_i, \ldots, Y_{\langle i+s-1\rangle}\right) \geq \frac{1}{s+1}\sum_{i=1}^{K} H\left(Y_i, \ldots, Y_{\langle i+s\rangle}\right),$$

*where we define $\langle i\rangle = i$ if $i \leq K$ and $\langle i\rangle = i - K$ if $i > K$.*

The resulting lower bounds, in their final form, are given in the following lemma.

**Lemma 5.** *Consider the multi-level, multi-user caching setup. Let $b \in \mathbb{N}^+$ and $t \in \{1, \ldots, K\}$. Furthermore, for every level $i$, let $s_i \in \mathbb{N}^+$ such that $s_i t \in \{d_i, \ldots, \lfloor K/2 \rfloor\}$. Then, for every memory $M$, the optimal rate can be bounded from below by:*

$$R^*(M) \geq \sum_{i=1}^{L} \lambda_i \cdot \min\left\{(s_i t - d_i + 1)U_i\,,\, \frac{N_i}{s_i b}\right\} - \frac{t}{b}M,$$

*where $\lambda_i$ is a constant (introduced for technical reasons) defined as: $\lambda_i = 1$ if $s_i t = d_i$; and $\lambda_i = \frac{1}{2}$ if $s_i t > d_i$.*

The full proof of Lemma 5 is given in Appendix A-A. However, we will here give a brief intuition behind the expression shown above. Every term in the sum corresponds to a level $i$. Consider a single level $i$. Roughly, if we ignore all the other levels in the summation, and assuming $d_i = 1$ for simplicity, then the inequality can be rearranged approximately as follows:

$$s_i b R + s_i t M \geq \min\left\{s_i t \cdot s_i b \cdot U_i\,,\, N_i\right\}.$$

This expression is essentially a cut-set bound, saying that, with $s_i t$ caches and $s_i b$ broadcast messages, up to $s_i t \cdot s_i b \cdot U_i$ files from level $i$ can be decoded (since there are $U_i$ users per cache for this level), unless this number excedes the total number of files $N_i$.

### *C. Approximate optimality*

In order to prove order-optimality of the memory-sharing scheme, we must use Lemma 5 with appropriate parameters. Our goal is to get a resulting lower bound on the optimal rate $R^*(M)$ such that the ratio between the achievable rate and $R^*(M)$—henceforth called the *gap*—is minimized. More specifically, the result we seek is a constant upper bound on the gap.

For technical reasons, several cases need to be considered for which different values are chosen for the parameters in Lemma 5. In particular, the values of choice in one case would violate the conditions imposed on them in another case. In

[9]For these inequalities to work, we have to take not just one cut-set bound per level, but an averaging of $K$ similar cut-set bounds that ends up including all caches an equal number of times.

this section, we will illustrate the proof methodology by approximating one of these cases. The complete and rigorous proof is given in Appendix A-B.

Recall the refined $M$-feasible partition introduced in Definition 1 and in (8). As previously mentioned, this partition only depends on the problem parameters, not on the achievability strategy. The cases of interest for the proof of order-optimality are:

- Case 1a: $I_1 = \emptyset$ and $J \neq \emptyset$;
- Case 1b: $I_1 = \emptyset$ and $J = \emptyset$;
- Case 2: $I_1 \neq \emptyset$.

There is also a special Case 0 for when $K$ is small (specifically, $K < D/\beta$). All the other cases assume $K$ is large.

The case that we will focus on in this section is Case 1a. Note that $I_1 = \emptyset \implies I = I_0 \cup I'$. For simplicity, we will assume that $d_i = 1$ for all $i$. Furthermore, most of the analysis will be approximate.

By Lemma 3, the achievable rate in Case 1a can be upper-bounded by:

$$R(M) \leq \sum_{h \in H} KU_h + \frac{2S_I^2}{M - T_J}, \tag{9}$$

since $S_I = \sum_{i \in I} \sqrt{N_i U_i}$ (see Definition 1). Now consider Lemma 5 with the following parameters:

$$\begin{aligned} & & t &= 1; \\ \forall h \in H, & & s_h &\approx \frac{1}{2} K; \\ \forall i \in I, & & s_i &\approx \frac{\sqrt{N_i/U_i}}{2\tilde{M}}; \\ \forall j \in J, & & s_j &= 1; \\ & & b &\approx 4\tilde{M}^2, \end{aligned}$$

where, as in Definition 1, $\tilde{M} = (M - T_J + V_I)/S_I \approx (M - T_J)/S_I$. The values of the $s_i$'s are very similar to those used in the single-level setup in [4]. Indeed, the levels with the smallest memory ($h \in H$) are handled using cut-set bounds that consider a fraction of the total number of caches, while the levels with the largest memory ($j \in J$) are handled with bounds that consider only one cache. Thus, the lemma lower-bounds the optimal rate by:

$$\begin{aligned} R^*(M) \geq & \sum_{h \in H} \min \left\{ s_h U_h, \frac{N_h}{s_h b} \right\} + \sum_{i \in I} \min \left\{ s_i U_i, \frac{N_i}{s_i b} \right\} \\ & + \sum_{j \in J} \min \left\{ s_j U_j, \frac{N_j}{s_j b} \right\} - \frac{M}{b}. \end{aligned}$$

Substituting the values of the parameters, and utilizing the inequalities on $\tilde{M}$ that define the refined $M$-feasible partition, we get:

$$\begin{aligned} R^*(M) &\geq \sum_{h \in H} \frac{1}{2} KU_h + \sum_{i \in I} \frac{\sqrt{N_i U_i} S_I}{2(M - T_J)} + \sum_{j \in J} \frac{N_j}{b} - \frac{M}{b} \\ &= \sum_{h \in H} \frac{1}{2} KU_h + \frac{S_I^2}{2(M - T_J)} - \frac{M - T_J}{b} \\ &= \sum_{h \in H} \frac{1}{2} KU_h + \frac{S_I^2}{2(M - T_J)} - \frac{(M - T_J) S_I^2}{4(M - T_J)^2} \\ &= \sum_{h \in H} \frac{1}{2} KU_h + \frac{S_I^2}{4(M - T_J)} \\ &\overset{(a)}{\geq} \frac{1}{8} \cdot R(M). \end{aligned}$$

The last inequality $(a)$ is due to (9).

Thus, the result is a constant multiplicative gap (8 in this example) between the achievable rate and the optimal rate for this regime. The full proof in Appendix A-B derives such gaps for each one of the four regimes mentioned above (Cases 0, 1a, 1b, and 2), and finally combines them into one multiplicative gap that holds for all cases.

## VI. The Single-User Setup

### A. Caching-and-delivery strategy: clustering

*Proof of Theorem 5:* Recall from Section IV-B how the memory is divided among the sets $H'$ and $I'$, defined in (6): all of the available memory is given to levels in $I'$, which is treated as one super-level. As a result, all requests for files from $H'$ must be handled by complete file transmissions from the BS. Since there are $\sum_{h\in H'} K_h$ users making such requests, the result is the same amount of transmissions in the worst case. Therefore, the message sent to all users requesting from a level in $H'$ has the following rate:

$$R_{H'} = \sum_{h\in H'} K_h. \tag{10}$$

For the set $I'$, now considered as one super-level, we use the single-level strategy from [5]. Although only a subset of the caches is active in our setup, the same strategy still applies. Indeed, the placement in [5] is a random sampling of the files in all the caches; we do the same placement in this case. In the delivery phase, we now know the caches to which the users of $I'$ connected. We perform a delivery as in [5], assuming that only these caches were ever present in the system.

For illustration, suppose that there were $K = 4$ caches. Furthermore, assume the partition $(H', I')$ was performed such that 3 users will request files from the set of levels $I'$. In the placement phase, we store a random sample of the files in $I'$ in each of the four caches. In the delivery phase, suppose the three users requesting files from $I'$ connect to caches 1, 2, and 4. In this case, the BS will send a complete file transmission for user 3 (since he requests a file from $H'$), and will treat users 1, 2, and 4 as though they were part of a single-level caching system with only three caches.

The rate required for $I'$ can be directly derived from Lemma 1, using $\sum_{i\in I'} K_i$ caches, $\sum_{i\in I'} N_i$ files, and 1 user per cache. In addition, we have, from (6), that $M \geq N_i/K_i$ for all $i \in I'$. This implies $M \geq (\sum_{i\in I'} N_i)/(\sum_{i\in I'} K_i)$, and hence the rate for $I'$ is:

$$R_{I'} = \max\left\{\frac{\sum_{i\in I'} N_i}{M} - 1\,,\, 0\right\}. \tag{11}$$

The maximization with zero is needed because it is possible to have $M > \sum_{i\in I'} N_i$.

By combining (10) with (11), we get a total broadcast rate of:

$$R(M) = R_{H'} + R_{I'} = \sum_{h\in H'} K_h + \max\left\{\frac{\sum_{i\in I'} N_i}{M} - 1, 0\right\}.$$

This proves Theorem 5. ∎

It will be helpful for the later analysis to refine the partition $(H', I')$ as follows.

**Definition 2.** *Define the following partition $(G, H, I, J)$ of the set of levels:*

$$\begin{aligned} G &= \{g : M < N_g/K_g \text{ and } K_g \leq 5 \text{ and } M \leq N_g/6\}\,; \\ H &= \{h : M < N_h/K_h \text{ and } K_h \geq 6\}\,; \\ I &= \{i : N_i/K_i \leq M \leq N_i/6\}\,; \\ J &= \{j : M > N_j/6\}\,. \end{aligned}$$

We rewrite the achievable rate in terms of this new partition:

$$R(M) \leq 5\cdot|G| + \sum_{h\in H} K_h + \frac{\sum_{i\in I} N_i}{M} + \left[\frac{\sum_{j\in J} N_j}{M} - 1\right]^+, \tag{12}$$

where $[x]^+ = \max\{x, 0\}$, and keeping in mind that $K_g \leq 5$ for $g \in G$. If we define $N_J = \sum_{j\in J} N_j$, we can upper-bound the last term by:

$$\left[\frac{N_J}{M} - 1\right]^+ \leq \begin{cases} N_J/M & \text{if } M < N_J/6; \\ 6\,(1 - M/N_J) & \text{if } N_J/6 \leq M < N_J; \\ 0 & \text{if } M \geq N_J. \end{cases} \tag{13}$$

Effectively, the set $J$ now behaves as a unit.

Much of what defines the above partition hinges on the following question: What happens if we know the user profile *a priori*? If we do, we could imagine a strategy of total separation of the levels: every level $i$ is treated as a single-level system with $K_i$ caches and users, and $N_i$ files. Each level can thus get the entirety of the memory of their specific $K_i$ caches. Using Lemma 1, we can calculate the total rate for this hypothetical situation as:

$$R'(M) = \sum_{i=1}^{L} \min\left\{K_i, \frac{N_i}{M}\right\}\left(1 - \frac{M}{N_i}\right)^+,$$

which can be approximated by:

$$R'(M) \leq 5 \cdot |G| + \sum_{h \in H} K_h + \sum_{i \in I} \frac{N_i}{M} + 6 \sum_{j \in J} \left(1 - \frac{M}{N_j}\right)^+.$$

Interestingly, the $G$, $H$, and $I$ terms in this rate expression are exactly the same as those in (12). Therefore, except for levels $j \in J$ where $M > N_j/6$, our scheme would not benefit from prior knowledge of the user profile. However, the set $J$ is in general limited by the lack of this knowledge. Indeed, if $M > N_j$ for all $j \in J$ but $M < N_J = \sum_{j \in J} N_j$, then it is possible to store any level in every cache, but it is not possible to store all levels in all caches. Thus knowing which level is at which cache can bring the rate down to zero.

Finally, $G$ and $H$, which are both subsets of $H'$ (though not necessarily a partition; the definition of $J$ allows some of its levels to be in $H'$ too) were separated because levels behave differently when their number of caches is very small. Specifically, a level in $H$ can transition into $I$ as the memory increases, but a level in $G$ immediately jumps to $J$ since $N_g/K_g > N_g/6$.

### *B. Outer bounds and approximate optimality*

As mentioned in Section IV-B, we use a cut-set bound to lower-bound the optimal rate. The idea is to send a certain number $b$ of broadcast messages $X_1, \ldots, X_b$ that serve certain requests. We choose these requests as follows. For every level $i \in G \cup H \cup I$, consider a certain number $s_i \leq K_i$ of caches. These caches are distinct across levels. For all the $b$ broadcasts, the users connected to these $s_i$ caches will altogether request $s_i b$ distinct files from level $i$ if there are that many; otherwise they request all $N_i$ files. For the levels in the set $J$, we collectively consider some $s_J$ caches (distinct from the rest). The users at these $s_J$ caches will use all $b$ broadcasts to decode as many files from the set $J$ as possible, up to $s_J b$ files. Let $n_J$ denote this number; it will be determined later.

If we let $S = \sum_{i \notin J} s_i + s_J$ be the total number of caches considered, then, by Fano's inequality:

$$\begin{aligned} bR + SM &\geq H\left(Z_1, \ldots, Z_S, X_1, \ldots, X_b\right) \\ &\geq \sum_{i \notin J} \min\left\{s_i b, N_i\right\} + n_J \\ R^*(M) &\geq \sum_{i \notin J} s_i \left(\min\left\{1, \frac{N_i}{s_i b}\right\} - \frac{M}{b}\right) \\ &\quad + s_J \left(\frac{n_J}{s_J b} - \frac{M}{b}\right) \\ &= \sum_{i \notin J} v_i + v_J. \end{aligned} \tag{14}$$

We will analyze each of the $v_i$ and $v_J$ terms separately. We identify two cases for which the analysis is slightly different.

The first case is when $M < 1/6$. Because of regularity condition (2), this implies $M < 1/6 < 1 \leq N_i/K_i$ for all levels $i$, and thus the achievable rate can be bounded by:

$$R(M) \leq \sum_{i=1}^{L} K_i. \tag{15}$$

The second case, which is more interesting, is when $M \geq 1/6$.

The details of the analysis are given in Appendix B. We will here give a brief outline of the procedure, with approximations to avoid burying the essence of the argument under technicalities. The idea is to lower-bound each term $v_i$ or $v_J$ to match the corresponding term in the rate expression. Let us focus on the case $M \geq 1/6$, and suppose we choose $b \approx 6M \geq 1$. For simplicity, we will only look at the sets $H$ and $I$.

Consider a level $h \in H$, and let $s_h \approx K_h/6$. Then, by the definition of $H$:

$$\frac{N_h}{s_h b} = \frac{N_h}{K_h M} \geq 1.$$

As a result,

$$v_h = \frac{K_h}{6}\left(1 - \frac{M}{6M}\right) = \frac{5}{36} \cdot K_h,$$

which matches the corresponding achievable rate term in (12), up to the constant $(5/36)$.

Now consider a level $i \in I$, and let $s_i \approx N_i/(6M)$. Then,

$$\frac{N_i}{s_i b} = \frac{N_i}{(N_i/6M) \cdot 6M} = 1,$$

and hence:

$$v_i = \frac{N_i}{6M}\left(1 - \frac{M}{6M}\right) = \frac{5}{36} \cdot \frac{N_i}{M}.$$

Again, this matches the corresponding achievable rate term in (12), up to the constant $(5/36)$.

Applying a similar procedure for every level, we get matching lower bounds (up to a constant) and thus prove Theorem 6.

## VII. Comparison of the Two Setups

In this section, we first compare the memory-sharing and the clustering strategies, and we explore the dichotomy among the two setups that is emphasized by the difference between strategies. We will then discuss why such a dichotomy exists, and explain the need for different lower bounds for each setup. Finally, we explore a new problem that combines both setups by including both multi-user and single-user levels.

### A. Comparing the two caching-and-delivery strategies

We have previously argued that memory-sharing is the best scheme to use in the multi-user case, while clustering is the near-optimal strategy in the single-user case. However, could one (or both) of these schemes be good enough for both situations? We will show, in this section, that it is not the case: in the single-user setup, memory-sharing can give a gap between its rate and the optimum that increases linearly with $L$; meanwhile, the rate achieved by clustering in the multi-user case can be arbitrarily far from the optimal rate. We give examples of these two cases.

Consider a multi-user setup with two levels. Suppose that there is enough memory so that both levels are to be partially stored in the caches. With the memory-sharing scheme, this means $I = \{1, 2\}$, which would, by Theorem 3, give a rate of approximately:

$$\begin{aligned} R &\approx \frac{\left(\sqrt{N_1 U_1} + \sqrt{N_2 U_2}\right)^2}{M} \\ &= \frac{1}{M}\left[N_1 U_1 + N_2 U_2 + 2\sqrt{N_1 U_1 N_2 U_2}\right]. \end{aligned}$$

On the other hand, if we had clustered the two levels into one, then this super-level would have $(N_1 + N_2)$ files and $(U_1 + U_2)$ users per cache, resulting in the following rate:

$$\begin{aligned} R &\approx \frac{(N_1 + N_2)(U_1 + U_2)}{M} \\ &= \frac{1}{M}\left[N_1 U_1 + N_2 U_2 + N_1 U_2 + N_2 U_1\right]. \end{aligned}$$

However, we know that the geometric mean of any two numbers is always smaller than their arithmetic mean. By considering the two numbers $N_1 U_2$ and $N_2 U_1$, we get:

$$2\sqrt{N_1 U_2 N_2 U_1} \leq N_1 U_2 + N_2 U_1,$$

and specifically the ratio between them can get arbitrarily large when the popularities of the two levels become significantly different (i.e., $U_2/N_2 \ll U_1/N_1$). Intuitively, if the two levels had similar popularities, then memory-sharing gives them similar amounts of memory, effectively merging them. However, if their popularities were very different, then they should be given drastically different portions of the memory.

Consider now the single-user case with $L$ levels, and suppose again that the memory is such that all levels will be partially stored. Let us assume that $N_1 = \cdots = N_L$. Using the clustering scheme, we get the following approximate rate:

$$R \approx \frac{N_1 + \cdots + N_L}{M} = \frac{L N_1}{M}.$$

However, with memory-sharing, we would get:

$$R \approx \frac{\left(\sqrt{N_1} + \cdots + \sqrt{N_L}\right)^2}{M} = \frac{L^2 N_1}{M},$$

which is larger by a factor of $L$. Essentially, we are sending $L$ broadcasts, one per level, when we could send just one broadcast for all $L$ levels.

### B. Analysis of the dichotomy between the setups

The dichotomy between the two extremes is striking. They require different strategies, and the strategy that is good for one setup is not so for the other. This suggests a fundamental difference between the two setups.

To understand this difference, consider what happens when sending a coded broadcast message. Each message targets a specific subset of users. If, in this subset, there exist two users that are connected to the same cache, then these users have access to the exact same side information. As a result, no coding can be done across these two users, and there is hence no benefit in including them in the same broadcast.

There are in fact two opposing forces at work on the caching-and-delivery strategy. The first is a popularity-centric force: it pushes on the strategy to allocate more memory to the more popular files. The second force is coding-centric: it encourages increasing the number of coding opportunities. These two forces are at odds, since coding *across levels* performs best when the files in these levels are given the same memory, regardless of popularity.

With that in mind, consider again Fig. 2 and Fig. 4. Notice how, in the multi-user setup, there are multiple rows of users, each of which consists of users from the *same* popularity level. Each such row is a complete set of users with no common caches: any additional users we add would have access to the same cache as some other user. Thus, it is sufficient to consider them in a broadcast transmission that is separate from all other rows. Since, as a result, no two levels will share the same broadcast message, it can only be beneficial to choose the best possible division of the memory, based on popularities. In the single-user setup, however, there is only one row of users that contains all the users from all the levels. It is hence possible to generate coding opportunities across levels. Merging is thus a better option in this situation, and merging is most efficient when all levels receive equal memory per file.

### C. The difference in the lower bounds

Complementing the difference in the achievable strategies between the two setups, we see a difference in the lower bounds to the optimal rate. In the multi-user case, we use a combination of cut-set bounds, one for each popularity level. However, we use a unified cut-set bound for all the levels in the single-user case. We elaborate on this difference in this section.

In the single-level setup studied in [4], cut-set bounds were given to lower-bound the optimal rate for every value of memory $M$. Depending on the value of $M$, a certain number of caches (and hence users) were considered and used in the cut-set bounds. Roughly speaking, about $M$ broadcast messages are sent to the users at $N/M$ distinct caches, allowing them to decode $M \cdot (N/M) = N$ files in total. In the multi-level setups (both multi-user and single-user), the lower bounds retain this idea. However, There are two crucial differences between the setups that force different choices of lower bounds.

The first difference is in the role of each cache vis-à-vis the popularity levels. In both setups, different levels are given different memory values. However, the *same* cache must be simultaneously used for *all* levels in the multi-user setup; in the single-user setup each cache is bound to a single level at any moment. As a result, a single cut-set bound can still encompass all levels in the single-user setup, but will not be enough in the multi-user setup.

The second difference is in the uncertainty of the user profile. In the single-user setup, there are situations where this uncertainty is significant enough to impact the achievable rate. This is especially true for levels nearing their maximal storage (the set $J$ in Definition 2), as described in Section VI-A. The lower bounds should incorporate this notion by considering demands from different levels at the same cache. As a result, a single cut-set bound unifying all levels (at least the levels involved in this uncertainty) becomes necessary.

### D. Mixing the setups

So far, we have looked at the two extremes: either all levels were represented at all the caches, or none of them were. A natural problem arises: that of studying intermediate cases. The simplest form such intermediate cases can take is one where levels of both types are present.

Specifically, there are two classes of popularity levels: $\mathcal{F}$ and $\mathcal{G}$. The class $\mathcal{F}$ consists of levels $i$ that are represented by exactly $U_i$ users *at every cache*. In contrast, there is exactly one row of users that represents all the levels in the class $\mathcal{G}$: each level $i \in \mathcal{G}$ is represented by $K_i$ of those users.

The most natural strategy to employ in this situation would be to superpose the multi-user and the single-user strategies. In particular, we divide the memory $M$ into $\gamma M$ and $(1-\gamma)M$, for some $\gamma \in [0,1]$. We give the first part to $\mathcal{F}$ and the second part to $\mathcal{G}$, and apply their respective strategies on their part of the memory. We believe this to be the best strategy, but proving its order-optimality requires developing new lower bounds that consider levels of both classes at the same time; this is part of our on-going work.

## VIII. Discussion and Numerical Evaluations

In the previous sections, we presented theoretical results for any given set of popularity levels and associated user access structures. However, in practice, what is available is a "continuous" popularity distribution over the entire set of $N$ files, and it is up to the designer to choose: (a) the number of popularity levels; (b) which files to assign to which level; and (c)

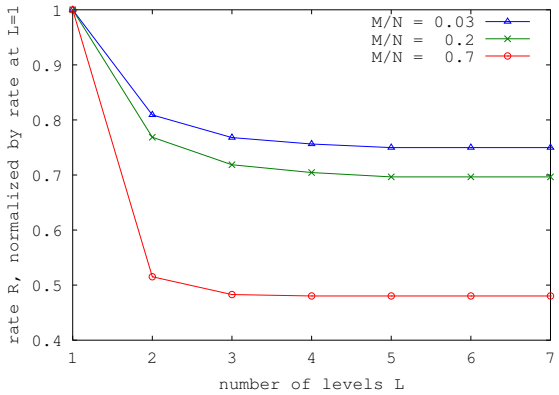


Fig. 11. Rate achieved by the memory-sharing scheme vs. number of levels, for different values of cache memory. For each $L$, we choose the $L$ levels that minimize the achievable rate (using brute-force search). For ease of comparison, the rate values have been normalized by the rate at $L = 1$.

the corresponding user access degree for each popularity level. For each such choice, our theoretical results characterize the minimum broadcast transmission rate, and we study, in this section, the impact of these choices on the transmission rate. Furthermore, while our theoretical model assumed that, for each popularity level, the number of users fixed, we relax this assumption here by allowing each user to randomly connect to one of the $K$ APs and request a file stochastically, according to the underlying popularity distribution. Finally, we will also compare the performance of our scheme with that of the traditional LFU approach, as well as the information-theoretic lower bounds presented before. These evaluations will only focus on the multi-user setup, as it is less discussed in the literature. Since the single-user setup utilizes a scheme similar to what is already in the literature for arbitrary distributions [8], [7], we do not feel it is necessary to include it in this discussion. However, we do provide a brief comparison of the clustering scheme with the literature in Section VIII-G

We use a YouTube dataset [37] for our evaluations. This dataset provides the number of views of videos in a set of $N \approx 500\,000$, over some period of time. Thus, these views can be thought of as approximating the popularity of the videos. Fig. 1 shows the popularity distribution of the videos (normalized number of views), which resembles a Zipf distribution similar to those commonly observed for multimedia content [17].

When using the YouTube dataset in the following sections, we will often omit the total number of users considered. This is because the total broadcast rate is always directly proportional to the total number of users, when the *fraction* of users per level is fixed. In this situation, the fraction is determined by the distribution in Fig. 1 and the levels into which the files were split, thus the total number of users will not affect the behavior of the system.

### A. Discretizing a continuous popularity distribution

Our first step is to divide the files in the YouTube dataset into a certain number of levels, based on the popularity profile in Fig. 1. Let there be $K = 75$ caches. We consider small, moderate, and large values of $M/N$ ($0.03$, $0.2$, and $0.7$) and set the user access degree $d_i = 1$ for every level $i$, so as to study the impact of the number of levels on the broadcast rate in isolation. For increasing values of $L$, we find the division of the files into $L$ levels that minimizes the rate achieved by the memory-sharing scheme using a brute-force search. We plot the minimum achievable rate versus $L$ in Fig. 11. As is easily apparent, while there is a significant gain in performance between treating all files as one level and dividing them into two levels, the gain decreases with diminishing returns as $L$ increases. This shows the importance of dividing files into multiple levels, but also suggests that 3–4 levels are sufficient to derive most of the benefits.

We remind the reader that the popularity profile in Fig. 1 is purely empirical. It is based on the number of views of the videos in the dataset, collected over some period of time. By the very nature of the data, if two videos have received, let's say, 1000 and 1100 views, then this does not really imply that the first file is more popular than the other. This is especially true of videos with very small number of views. By grouping files into popularity classes, we acknowledge the difference in popularity of very different files, while simultaneously not distinguishing between files whose *empirical* popularities are close.

### B. Impact of multi-access on the achievable rate

To study the effect of multi-access in isolation, we will fix the partition we use to divide the files into different levels and then look at different multi-access structures. Suppose again $K = 75$, and consider $L = 2$ levels, with $N_1 = 0.2N$, and $N_2 = 0.8N$ files. We plot in Fig. 12 the broadcast rate of our scheme as a function of the normalized memory $M/N$, for four different access structures $(d_1, d_2)$: $(1, 1)$, $(1, 2)$, $(2, 1)$, and $(2, 2)$.

As one would expect, allowing for multi-access greatly improves the transmission rate. For example, the rate for the multi-access system with $(d_1 = 2, d_2 = 2)$ is smaller than the rate for the single-access system with $(d_1 = 1, d_2 = 1)$. The cases $(d_1 = 1, d_2 = 2)$ and $(d_1 = 2, d_2 = 1)$ provide a more interesting comparison. For small memory size $M$, the former gives a

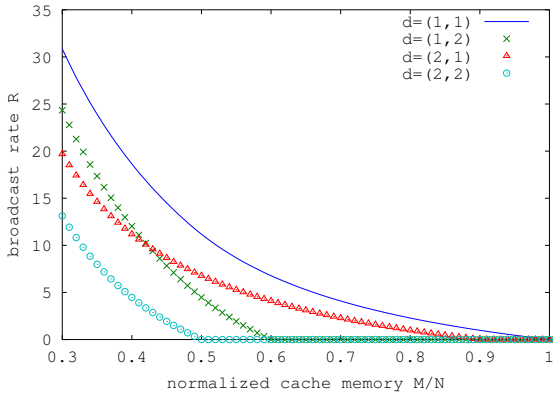


Fig. 12. Achievable rate vs. cache memory in a two-level setup, for different access structures.

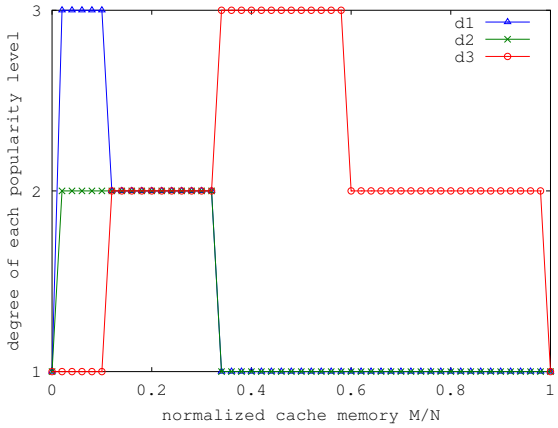


Fig. 13. Optimal access structure vs. memory, with $d_{\max} = 3$, $d_{\text{avg}} = 2$.

lower rate since the cache memory mainly contains files from level 1, and so giving higher access to level 1 is more beneficial in reducing the rate. On the other hand, as $M$ grows and files from level 2 start occupying a significant portion of the memory, it becomes more efficient to give higher access to level 2 since it has many more files than level 1.

While greater cache access helps reduce the rate, there is also a cost associated with it in terms of the increased delay in gathering data from multiple APs, as well as a reduced rate as a user connects to farther APs. In general, for a given multi-level setup with parameters $L$, $K$, $\{N_i, U_i\}$, and $M$, such a cost can be included in the rate optimization framework as one or more inequalities of the form $\text{cost}_j(K, \{U_i, d_i\}_i) \leq C_j$, for some maximum cost $C_j$. The above optimization problem can be numerically solved by a designer in order to identify the optimal access structure for the multi-level system under consideration. However, to derive some intuition about how the costs impact the optimal multi-access structure, let us consider a setup with $L = 3$ levels, and with $N_1 = 0.04N$, $N_2 = 0.13N$, and $N_3 = 0.83N$ files in the three levels. Say we want to include both a maximum degree constraint $d_i \leq 3$ for each level $i$, as well as an average degree constraint $(\sum_i U_i d_i)/U \leq 2$. Then, Fig. 13 plots the optimal access structure vs. the normalized memory size. As before, when the memory is small, the optimal access structure is one which satisfies $d_1 \geq d_2 \geq d_3$, but this relation becomes reversed as the memory increases.

### C. Stochastic variations in user profiles

The theoretical setup and results presented in the previous sections assumed a symmetric and deterministic user profile across all the APs. In particular, exactly $U_i$ users are assigned to each AP to request files from level $i$. This section aims at evaluating the robustness of memory-sharing to asymmetry and stochasticity in the user profiles across caches.

We consider a setup where each of the $KU$ users in the system randomly connects to one of the $K$ APs and requests a file stochastically, according to the YouTube popularity distribution in Fig. 1. The scheme we use here is a simple variation of the one for the worst-case setup: the files are split into two levels, and the placement is done using memory-sharing based on their average popularity. The delivery phase is almost identical to the worst-case delivery, with the exception that, because of a lack of determinism in the user profile, not all caches will have the same number of users per level. This is handled by simply trying to group as many users of the same level as possible in each broadcast transmission.

We ran simulations for this setup using the above strategy, and we plot here the empirically achieved rate against the cache memory in Fig. 14. For comparison, we also show the rate predicted by our theoretical model, which splits the files into two

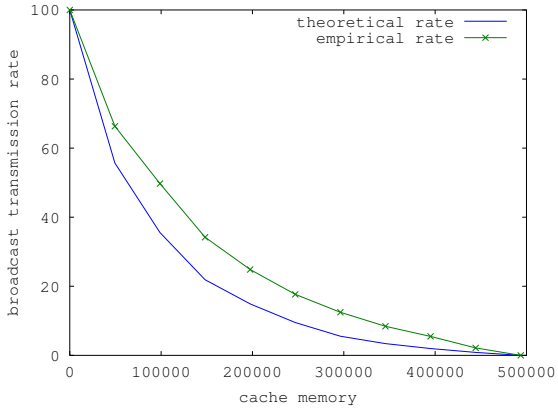


Fig. 14. Comparison of the theoretical rate with the empirical rate, based on simulations of demands over the YouTube dataset, with 5 caches and 100 total users. The theoretical rate is off by a factor of up to 2.8 from the empirical.

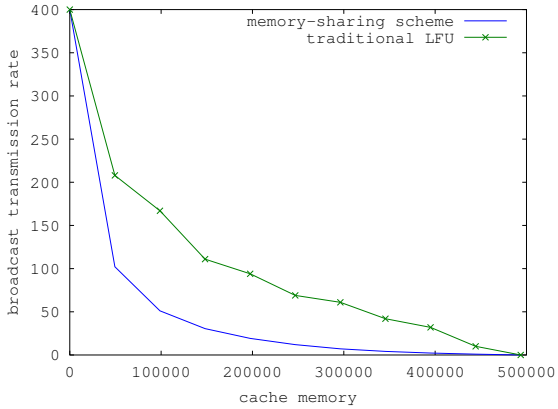


Fig. 15. Comparison of the memory-sharing scheme with traditional LFU. The memory-sharing scheme achieves up to a factor-14.5 in gain over LFU.

levels and assumes a symmetric user profile across the caches. Clearly, the theory very closely predicts the empirical results for a random user profile, thus demonstrating the robustness of our theoretical results to stochastic variations across APs and justifying their utility in practice.

### *D. Comparison with Least-Frequently Used (LFU)*

In this section, we compare the performance of memory-sharing with that of the traditional LFU scheme using simulations on the YouTube data. For any memory size $M$, LFU fully stores the $M$ most popular files, so that requests for more popular files are completely served from the cache, and requests for less popular files are fully handled by the BS transmission. The results, given in Fig. 15, show the superiority of memory-sharing.

### *E. Comparison with different memory-sharing strategies*

The memory-sharing scheme proposed for the multi-user setup relies on a very specific division of the memory among the levels. A natural question that arises is if this memory-sharing is necessary. In this section, we compare its behavior with that of different, perhaps more natural, memory-sharing schemes. In particular, we compare with:

- **Coded LFU:** This scheme gives as much memory as possible for the most popular levels. We call it "coded LFU" because it allocates memory in a similar way to traditional LFU, but delivers a coded broadcast message to the users.
- **Uniform memory-sharing:** This scheme ignores all popularities and gives the same amount of memory to all files.

Our memory-sharing scheme—which we will call "optimal memory-sharing" in this section to distinguish it from the others—performs better than the other two. We note, however, that coded LFU and uniform memory-sharing are both special cases of optimal memory-sharing. Indeed, if all levels have exactly the same popularities, then optimal memory-sharing will give all files equal memory. On the other hand, when the level popularities are separted enough, then optimal memory-sharing will end up always prioritizing the most popular levels, and thus it reduces to coded LFU. In the middle case, it provides benefits that both other schemes lack.

As an example, consider the following setup. There are $K = 100$ caches, and $L = 3$ levels with $(N_1, N_2, N_3) = (2000, 5000, 50\,000)$. We have $(U_1, U_2, U_3) = (20, 10, 5)$, and single-access degree $(d_1, d_2, d_3) = (1, 1, 1)$. Then, optimal memory-sharing performs more than 6 times better than both other strategies.

### F. Numerical gap

As discussed in Section IV, the multiplicative gap in Theorem 4 results from many generous approximations in bounding the achievable rate. Numerical results suggest that this gap is in fact much smaller. We here give a few examples of these results.

- If $K = 10$, $L = 3$, $(N_1, N_2, N_3) = (500, 1500, 8000)$, $(U_1, U_2, U_3) = (9, 5, 1)$, and $(d_1, d_2, d_3) = (1, 3, 5)$, then we get a gap of approximately 6.
- Increasing the number of caches and files: if $K = 20$, $L = 3$, $(N_1, N_2, N_3) = (200, 20\,000, 800\,000)$, $(U_1, U_2, U_3) = (10, 5, 1)$, and $(d_1, d_2, d_3) = (1, 1, 1)$, then we get a gap of approximately 6.8.
- The same setup as the previous point, with access degrees of $(d_1, d_2, d_3) = (1, 2, 3)$ gives a gap of about 7.6.

In most examples we have tried, the gap was in the range 5-10, regardless of the access degree. In the worst case, the largest gap our numerics have shown was about 45, in a situation with $L = 3$ and $D = 5$.

### G. Relation of the single-user setup to the literature

As previously mentioned, the clustering scheme we proposed for the single-user setup is similar to strategies discussed in the literature for similar setups with stochastic demands [8], [7]. In these setups, users request files based on a probability distribution, and the average rate is analyzed (as opposed to the worst-case rate in our case). In both cases, the proposed strategy is to divide the files into two sets, based on some threshold, and store only the most popular files.

Recall that our strategy clusters the levels $i$ whose popularity is such that $M \geq N_i/K_i$. The popularity $p_i$ of a file of level $i$ is proportional to $K_i/N_i$, and hence if we normalize so that the sum of popularities is 1, we get:

$$p_i = \frac{K_i/N_i}{\sum_{j=1}^{L} N_j \cdot K_j/N_j} = \frac{K_i}{KN_i}.$$

Thus, the condition $M \geq N_i/K_i$ can be rewritten as $p_i \geq 1/KM$. This is exactly the threshold used in [8] to determine which files to cluster and store in the caches, and which files to leave out of the caches.

## REFERENCES


[1] "Cisco visual networking index (VNI) global mobile data traffic forecast update," 2013, http://www.gsma.com/spectrum/wp-content/uploads/2013/03/Cisco_VNI-global-mobile-data-traffic-forecast-update.pdf.

[2] "Qualcomm: The 1000x challenge," 2013, http://www.qualcomm.com/solutions/wireless-networks/technologies/1000x-data/small-cells.

[3] M. R. Korupolu, C. G. Plaxton, and R. Rajaraman, "Placement algorithms for hierarchical cooperative caching," in *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1999, pp. 586–595.

[4] M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *IEEE Transactions on Information Theory*, vol. 60, no. 5, pp. 2856–2867, May 2014.

[5] ——, "Decentralized coded caching attains order-optimal memory-rate tradeoff," *Networking, IEEE/ACM Transactions on*, vol. 23, no. 4, pp. 1029–1040, Aug 2015.

[6] U. Niesen and M. A. Maddah-Ali, "Coded caching with nonuniform demands," in *Computer Communications Workshops (INFOCOM WKSHPS), 2014 IEEE Conference on*, Apr. 2014, pp. 221–226.

[7] M. Ji, A. M. Tulino, J. Llorca, and G. Caire, "On the average performance of caching and coded multicasting with random demands," in *IEEE International Symposium on Wireless Communication Systems*, Aug. 2014.

[8] J. Zhang, X. Lin, and X. Wang, "Coded caching under arbitrary popularity distributions," in *Proc. Information Theory and Applications Workshop (ITA)*, Feb. 2015.

[9] J. Hachem, N. Karamchandani, and S. Diggavi, "Multi-level coded caching," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Jun. 2014.

[10] ——, "Multi-level coded caching over heterogeneous wireless networks," in *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, 2015.

[11] ——, "Effect of number of users in multi-level coded caching," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Jun. 2015.

[12] D. Wessels, *Web Caching*, N. Torkington, Ed. Sebastopol, CA, USA: O'Reilly & Associates, Inc., 2001.

[13] S. Borst, V. Gupta, and A. Walid, "Distributed caching algorithms for content distribution networks," in *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, 2010, pp. 1478–1486.

[14] B. Tan and L. Massoulié, "Optimal content placement for peer-to-peer video-on-demand systems," *IEEE/ACM Transactions on Networking*, vol. 21, no. 2, pp. 566–579, Apr. 2013.

[15] J. Llorca, A. M. Tulino, K. Guan, J. Esteban, M. Varvello, N. Choi, and D. C. Kilper, "Dynamic in-network caching for energy efficient content delivery," in *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, 2013, pp. 245–249.

[16] A. Wolman, G. M. Voelker, N. Sharma, N. Cardwell, A. Karlin, and H. M. Levy, "On the scale and performance of cooperative web proxy caching," in *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP)*, 1999, pp. 16–31.

[17] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and zipf-like distributions: evidence and implications," in *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, 1999, pp. 126–134.

[18] D. Applegate, A. Archer, V. Gopalakrishnan, S. Lee, and K. K. Ramakrishnan, "Optimal content placement for a large-scale VoD system," in *Proceedings of the ACM International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, 2010, pp. 4:1–4:12.

[19] N. Golrezaei, K. Shanmugam, A. G. Dimakis, A. F. Molisch, and G. Caire, "Femtocaching: Wireless video content delivery through distributed caching helpers." in *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, Mar. 2012, pp. 1107–1115.

[20] R. Pedarsani, M. A. Maddah-Ali, and U. Niesen, "Online coded caching," *pre-print*, 2013, http://arxiv.org/abs/1311.3646.

[21] S. Wang, W. Li, X. Tian, and H. Liu, "Fundamental limits of heterogenous cache," *CoRR*, vol. abs/1504.01123, 2015. [Online]. Available: http://arxiv.org/abs/1504.01123

[22] J. Zhang, X. Lin, C.-C. Wang, and X. Wang, "Coded caching for files with distinct file sizes," in *Information Theory (ISIT), 2015 IEEE International Symposium on*, June 2015, pp. 1686–1690.
[23] H. Ghasemi and A. Ramamoorthy, "Improved lower bounds for coded caching," in *Information Theory (ISIT), 2015 IEEE International Symposium on*, June 2015, pp. 1696–1700.
[24] A. Sengupta, R. Tandon, and T. Clancy, "Improved approximation of storage-rate tradeoff for caching via new outer bounds," in *Information Theory (ISIT), 2015 IEEE International Symposium on*, June 2015, pp. 1691–1695.
[25] G. Vettigli, M. Ji, A. M. Tulino, J. Llorca, and P. Festa, "An efficient coded multicasting scheme preserving the multiplicative caching gain," *CoRR*, vol. abs/1511.07531, 2015. [Online]. Available: http://arxiv.org/abs/1511.07531
[26] K. Shanmugam, M. Ji, A. M. Tulino, J. Llorca, and A. G. Dimakis, "Finite length analysis of caching-aided coded multicasting," *CoRR*, vol. abs/1508.05175, 2015. [Online]. Available: http://arxiv.org/abs/1508.05175
[27] N. Karamchandani, U. Niesen, M. A. Maddah-Ali, and S. Diggavi, "Hierarchical coded caching," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Jun. 2014.
[28] ——, "Hierarchical coded caching," *CoRR*, vol. abs/1403.7007, 2014. [Online]. Available: http://arxiv.org/abs/1403.7007
[29] M. Ji, G. Caire, and A. Molisch, "Wireless device-to-device caching networks: Basic principles and system performance," *Selected Areas in Communications, IEEE Journal on*, vol. 34, no. 1, pp. 176–189, Jan 2016.
[30] S. Jeon, S. Hong, M. Ji, G. Caire, and A. F. Molisch, "Wireless multihop device-to-device caching networks," *CoRR*, vol. abs/1511.02574, 2015. [Online]. Available: http://arxiv.org/abs/1511.02574
[31] S. P. Shariatpanahi, A. S. Motahari, and B. H. Khalaj, "Multi-server coded caching," *CoRR*, vol. abs/1503.00265, 2015. [Online]. Available: http://arxiv.org/abs/1503.00265
[32] M. Ji, A. M. Tulino, J. Llorca, and G. Caire, "Caching-aided coded multicasting with multiple random requests," *CoRR*, vol. abs/1511.07542, 2015. [Online]. Available: http://arxiv.org/abs/1511.07542
[33] S. Gitzenis, G. S. Paschos, and L. Tassiulas, "Asymptotic laws for joint content replication and delivery in wireless networks," *IEEE Transactions on Information Theory*, vol. 59, no. 5, pp. 2760–2776, May 2013.
[34] S. Ioannidis, L. Massoulie, and A. Chaintreau, "Distributed caching over heterogeneous mobile networks," in *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 2010, pp. 311–322.
[35] E. Altman, K. Avrachenkov, and J. Goseling, "Coding for caches in the plane," *pre-print*, 2013, http://arxiv.org/abs/1309.0604.
[36] J. Y. Yang and B. Hajek, "Single video performance analysis for video-on-demand systems," *pre-print*, 2013, http://arxiv.org/abs/1307.0849.
[37] "UMass trace repository," http://traces.cs.umass.edu/index.php/network/network.
[38] J. Jiang, N. Marukala, and T. Liu, "Symmetrical multilevel diversity coding and subset entropy inequalities," *IEEE Transactions on Information Theory*, vol. 60, no. 1, pp. 84–103, 2014.

## APPENDIX A
## PROOFS FOR THE MULTI-USER SETUP

### A. Proof of Lemma 5

*1) A small example for illustration:* Before we give the full proof of Lemma 5, we will start with a simple example for illustration. Consider a multi-level, multi-user caching system with $K = 6$ caches and $L = 3$ levels. Suppose that $U_1 = U_2 = U_3 = 1$, and let $N_1$, $N_2$, and $N_3$ be some large numbers (their exact value is not important for this example). Finally, assume a single-access structure for all users, i.e., $d_1 = d_2 = d_3 = 1$.

As we have discussed in Section V-B, the lower bounds on the optimal rate that we wish to obtain are a sum of $L$ cut-set bounds, each pertaining to a single level. A cut-set bound for level $i$ consists of a certain number of caches and broadcast messages, such that the users at these caches can use the broadcast to cooperatively decode a set of files from level $i$. For example, if we consider the $sU_i$ users of level $i$ connected to some $s$ caches, we can send $b$ broadcast messages, tailored for the correct request vectors, so that the users can collectively decode $\min\{sU_i \cdot b, N_i\}$ files of level $i$. Furthermore, the values of $s$ and $b$ are usually chosen so that the final bound matches the achieved individual rate of level $i$. Thus, they depend on value of the rate and hence on the memory available to the level.

In our example, suppose that the memory we expect each level to receive dictates the following number of caches to consider: a single cache for level 1; two caches for level 2; and three caches for level 3. We define $s_1 = 1$, $s_2 = 2$, and $s_3 = 3$ to be these numbers. In order to group together the cut-set bounds of all the levels with their different numbers of caches, we resort to the sliding-window subset entropy inequality presented in Lemma 4.

In this example, we will consider six separate broadcast messages. To avoid burying the essence of the argument under technicalities, we omit the request vectors from the discussion of the example. For simplicity, we will assume that the messages are chosen so that, any time we encounter a group of $p$ users and $q$ messages, the users are able to decode a total of $pq$ messages. This relaxation will be discarded when we discuss the general proof.

Let the caches be labeled with $Z_1, \ldots, Z_6$ and the broadcast messages with $X_1, \ldots, X_6$. We start with the level with the smallest $s_i$, in this case level 1. Since $s_1 = 1$, this means its cut-set bound will consider a single cache. In fact, the cut-set bound should have the following form by Fano's inequality:

$$H(Z_1, X_1) \geq H(Z_1, X_1|\mathcal{W}^1) + 1 \cdot F,$$

where $\mathcal{W}^1$ represents the set of files of level 1 that got decoded. In this case, the single level-1 user at cache $Z_1$ can only decode one file when given just $X_1$, hence the $1 \cdot F$ term. However, to take full advantage of the sliding-window entropy

inequality, we will take the average over six cut-set bounds for the same level, one for each cache:

$$\begin{aligned} RF + MF &\geq \frac{1}{6}\left[H(Z_1, X_1) + \cdots + H(Z_6, X_6)\right] \\ &\geq \frac{1}{6}\left[H(Z_1, X_1|\mathcal{W}^1) + \cdots + H(Z_6, X_6|\mathcal{W}^1)\right] \\ &\quad + 1 \cdot F. \end{aligned}$$

After completing the cut-set bound for level 1, we now attempt to transition to the cut-set bound for level 2. Because we have taking the average of six instances of cut-set bounds, we can use Lemma 4 to obtain 6 new cut-set bounds with $s_2 = 2$ caches each (for ease of notation, we write $Y_k = (Z_k, X_k)$):

$$\begin{aligned} RF + MF \;\geq\; & \frac{1}{6}\left[H(Y_1|\mathcal{W}^1) + H(Y_2|\mathcal{W}^1) + H(Y_3|\mathcal{W}^1)\right. \\ & \left. + H(Y_4|\mathcal{W}^1) + H(Y_5|\mathcal{W}^1) + H(Y_6|\mathcal{W}^1)\right] \\ & + 1 \cdot F \\ \overset{(a)}{\geq}\; & \frac{1}{6} \cdot \frac{1}{2}\left[H(Y_1, Y_2|\mathcal{W}^1) + H(Y_2, Y_3|\mathcal{W}^1)\right. \\ & + H(Y_3, Y_4|\mathcal{W}^1) + H(Y_4, Y_5|\mathcal{W}^1) \\ & \left. + H(Y_5, Y_6|\mathcal{W}^1) + H(Y_6, Y_1|\mathcal{W}^1)\right] \\ & + 1 \cdot F \\ \overset{(b)}{\geq}\; & \frac{1}{6} \cdot \frac{1}{2}\left[H(Y_1, Y_2|\mathcal{W}^1, \mathcal{W}^2) + H(Y_2, Y_3|\mathcal{W}^1, \mathcal{W}^2)\right. \\ & + H(Y_3, Y_4|\mathcal{W}^1, \mathcal{W}^2) + H(Y_4, Y_5|\mathcal{W}^1, \mathcal{W}^2) \\ & \left. + H(Y_5, Y_6|\mathcal{W}^1, \mathcal{W}^2) + H(Y_6, Y_1|\mathcal{W}^1, \mathcal{W}^2)\right] \\ & + 1 \cdot F + 4 \cdot F. \end{aligned}$$

Here, inequality $(a)$ uses Lemma 4, while inequality $(b)$ uses Fano's inequality on the level-2 cut-set bounds. Since each bound involves two users and two broadcast messages, the total number of files decoded is 4, hence the $4 \cdot F$ term. The set of decoded level-2 files is denoted by $\mathcal{W}^2$.

We proceed again with the transition from level 2 to level 3. Just like before, we first apply Lemma 4 to obtain entropy terms with the correct number of caches $s_3$, and then apply Fano's inequality to decode files from level 3, labeled $\mathcal{W}^3$.

$$\begin{aligned} RF + MF \;\geq\; & \frac{1}{6} \cdot \frac{1}{2}\left[H(Y_1, Y_2|\mathcal{W}^1, \mathcal{W}^2) + H(Y_2, Y_3|\mathcal{W}^1, \mathcal{W}^2)\right. \\ & + H(Y_3, Y_4|\mathcal{W}^1, \mathcal{W}^2) + H(Y_4, Y_5|\mathcal{W}^1, \mathcal{W}^2) \\ & \left. + H(Y_5, Y_6|\mathcal{W}^1, \mathcal{W}^2) + H(Y_6, Y_1|\mathcal{W}^1, \mathcal{W}^2)\right] \\ & + 1 \cdot F + 4 \cdot F \\ \overset{(a)}{\geq}\; & \frac{1}{6} \cdot \frac{1}{3}\left[H(Y_1, Y_2, Y_3|\mathcal{W}^1, \mathcal{W}^2)\right. \\ & + H(Y_2, Y_3, Y_4|\mathcal{W}^1, \mathcal{W}^2) \\ & + \cdots \\ & \left. + H(Y_6, Y_1, Y_2|\mathcal{W}^1, \mathcal{W}^2)\right] \\ & + 1 \cdot F + 4 \cdot F \\ \overset{(b)}{\geq}\; & \frac{1}{6} \cdot \frac{1}{3}\left[H(Y_1, Y_2, Y_3|\mathcal{W}^1, \mathcal{W}^2, \mathcal{W}^3)\right. \\ & + H(Y_2, Y_3, Y_4|\mathcal{W}^1, \mathcal{W}^2, \mathcal{W}^3) \\ & + \cdots \\ & \left. + H(Y_6, Y_1, Y_2|\mathcal{W}^1, \mathcal{W}^2, \mathcal{W}^3)\right] \\ & + 1 \cdot F + 4 \cdot F + 9 \cdot F. \end{aligned}$$

Again, inequality $(a)$ uses Lemma 4, and inequality $(b)$ uses Fano's inequality on level 3, which is considering groups of 3 users and 3 broadcast messages, allowing the decoding of a total of 9 files.

Finally, by the non-negativity of entropy, the lower bound becomes:

$$\begin{aligned} RF + MF &\geq (1 + 4 + 9) \cdot F \\ R + M &\geq 14. \end{aligned}$$

*2) The general proof:* The process involved in the general case is similar to the one shown in the example above. We start with the level with the smallest $s_i$ and write an average of $K$ cut-set bounds for this level, one for each sequence of consecutive $s_i$ caches. After applying Fano's inequality, we use Lemma 4 to transition to the level with the next smallest $s_i$. In this discussion, the popularity of the levels is not important, but the order of their $s_i$'s is. Specifically, if $s_i < s_j$, it does not matter which of $i$ and $j$ is more popular. Thus, we can assume without loss of generality that $s_1 \leq \cdots \leq s_L$.

In the example, we started with cut-set bounds with a single cache each. For technical reasons, this will not give us good enough bounds in general, and so the initial cut-set bounds will consist of $t$ caches, for a general $t \in \{1, \ldots, K\}$. In fact, every consecutive $t$ caches will be clustered into an inseperable group. The group consisting of the $t$ caches that start with cache $k$ is labeled as:

$$\mathcal{Z}_k^t = \left(Z_k, \ldots, Z_{\langle k+t-1 \rangle}\right),$$

where $\langle m \rangle$ is defined for integers $m$ as in Lemma 4, i.e., $\langle m \rangle = m$ if $m \leq K$ and $\langle m \rangle = m - K$ if $m > K$. To every cache group $\mathcal{Z}_k^t$, we associate a broadcast-message group $\mathcal{X}_k^b$, which consists of $b$ messages serving different user demands.

Recall that a level-$i$ user needs to connect to $d_i$ *consecutive* caches in order to decode whichever file he has requested. Applying Lemma 4 should hence keep only consecutive caches in the same cut-set bounds, which will allow for the maximum number of users to be active in the decoding of the files (and thus produce a larger number of files from the same cut-set bound). While this was fairly simple to ensure in the example above, we must show that we can still do it even after introducing the $t$-groups of caches. This is done in the next paragraph.

Let $g$ be the GCD of $t$ and $K$. Then, the following sequence:

$$\left(\mathcal{Z}_1^t, \mathcal{Z}_{\langle t+1 \rangle}^t, \ldots, \mathcal{Z}_{\langle (K/g)t+1 \rangle}^t\right),$$

starts at cache 1 and ends at cache $K$. Each cache $Z_k$ appears exactly $(t/g)$ times in the sequence. Furthermore, every pair of caches that are consecutive in the sequence are also consecutive in the system. For example, suppose $t = 4$ and $K = 6$. Then their GCD would be $g = 2$, with $K/g = 3$, and the sequence would be:

$$\Big((Z_1, Z_2, Z_3, Z_4), (Z_5, Z_6, Z_1, Z_2), (Z_3, Z_4, Z_5, Z_6)\Big).$$

Notice that every cache appears in the sequence $t/g = 2$ times, and that consecutive caches remain so in the sequence.

We are now ready to prove the lower bounds. Start with an average of $K/g$ cut-set bounds consisting of one group of $t$ caches (and their associated broadcast messages) each:

$$bRF + tMF \geq \frac{g}{K} \sum_{k=1}^{K/g} H\left(\mathcal{Z}_{(k-1)g+1}^t, \mathcal{X}_{(k-1)g+1}^b\right).$$

This follows from $bR + tM \geq H(\mathcal{Z}_k^t, \mathcal{X}_k^b)$ for all $k$. Notice how all the indices $[(k-1)g + 1]$ are the same modulo $g$. We can include all other caches to get:

$$bRF + tMF \geq \frac{1}{K} \sum_{k=1}^{K} H\left(\mathcal{Z}_k^t, \mathcal{X}_k^b\right).$$

The first step is to use Lemma 4 to obtain cut-set bounds with $s_1 t$ caches for level 1.

$$\begin{aligned} bRF + tMF &\geq \frac{1}{K} \sum_{k=1}^{1} H\left(\mathcal{Z}_k^t, \mathcal{X}_k^b\right) \\ &\geq \frac{1}{K} \cdot \frac{1}{s_1} \sum_{k=1}^{K} H\Big(\mathcal{Z}_k^t, \mathcal{Z}_{\langle k+t \rangle}^t, \ldots, \mathcal{Z}_{\langle k+(s_1-1)t \rangle}^t, \\ &\qquad\qquad \mathcal{X}_k^b, \mathcal{X}_{\langle k+t \rangle}^b, \ldots, \mathcal{X}_{\langle k+(s_1-1)t \rangle}^b\Big) \end{aligned}$$

Each entropy term in the sum now consists of $s_1 t$ *consecutive* caches and $s_1 b$ broadcast messages. For simplicity, we write it as $\hat{H}_k(s_1 t, s_1 b)$, for $k \in \{1, \ldots, K\}$, and its conditional version as $\hat{H}_k(s_1 t, s_1 b | Q)$ for any random variable $Q$. Let $\mathcal{W}^1$ denote the set of level-1 files that can be decoded in every cut-set bound in the sum, and let $p_1$ be its size, to be determined later. For technical reasons, we will only apply Fano's inequality on a fraction of the cut-set bounds; the rest are simply lower-bounded by conditioning. The meaning and value of this fraction is discussed below, but for now we will call it $\lambda_1$. Then, we use Fano's inequality to get:

$$\begin{aligned} bRF + tMF &\geq \frac{1}{K} \cdot \frac{1}{s_1} \sum_{k=1}^{K} \hat{H}_k(s_1 t, s_1 b) \\ &\geq \frac{1}{K} \cdot \frac{1}{s_1} \sum_{k=1}^{K} \hat{H}_k(s_1 t, s_1 b | \mathcal{W}^1) + \frac{\lambda_1 p_1}{s_1} \cdot F. \end{aligned}$$

By applying Lemma 4 again, we obtain cut-set bounds pertaining to level 2. We lower-bound a fraction $\lambda_2$ of them in turn using Fano's inequality, and repeat for all levels. Thus, we have:

$$\begin{aligned} bRF + tMF &\geq \frac{1}{K} \cdot \frac{1}{s_1} \sum_{k=1}^{K} \hat{H}_k(s_1 t, s_1 b | \mathcal{W}^1) + \frac{\lambda_1 p_1}{s_1} \cdot F \\ &\overset{(a)}{\geq} \frac{1}{K} \cdot \frac{1}{s_2} \sum_{k=1}^{K} \hat{H}_k(s_2 t, s_2 b | \mathcal{W}^1) + \frac{\lambda_1 p_1}{s_1} \cdot F \\ &\overset{(b)}{\geq} \frac{1}{K} \cdot \frac{1}{s_2} \sum_{k=1}^{K} \hat{H}_k(s_2 t, s_2 b | \mathcal{W}^1, \mathcal{W}^2) \\ &\quad + \frac{\lambda_1 p_1}{s_1} \cdot F + \frac{\lambda_2 p_2}{s_2} \cdot F \\ &\geq \cdots \\ &\geq \sum_{i=1}^{L} \frac{\lambda_i p_i}{s_i} \cdot F. \end{aligned} \tag{16}$$

where the inequality marked with $(a)$ uses Lemma 4 and the one marked with $(b)$ uses Fano's inequality.

For any $m < K$ *consecutive* caches, the number of level-$i$ users that are connected to $d_i$ of those caches is exactly $(m - d_i + 1)U_i$. Therefore, given $s_i t < K$ consecutive caches and $s_i b$ broadcast messages, the users at these caches should, in principle, decode up to $p_i = \min\{(s_i t - d_i + 1)U_i \cdot s_i b, N_i\}$ files from level $i$. This requires choosing the broadcast messages for the correct user demands. However, consider a pair of broadcast message $X'$ and user $u$ that appears in multiple cut-set bounds. In other words, the user $u$ is expected to decode a file using the message $X'$ on multiple occasions. We must ensure that there are no contradictions, i.e., that the message $X'$ always delivers the same file to user $u$.

It turns out it is not always possible to do that and still be able to decode the maximal number of files at every cut-set bound. However, because $s_i t \leq K/2$ (as constrained in the statement of Lemma 5), we can get around this issue by "ignoring" half of the cut-set bounds. To understand this, suppose we wish to determine the user demand vector of each broadcast message to be sent. Starting with $\hat{H}_1(s_i t, s_i b)$: there are a total of $(s_i t - d_i + 1)U_i$ relevant users, and $s_i b$ broadcasts. We can design these broadcasts to allow these users to decode all $p_i$ files. Next, we move on to $\hat{H}_2(s_i t, s_i b)$. In this entropy term, exactly $tU_i$ users from the previous term are replaced by $tU_i$ *brand new* users; the other users are still the same. Thus, if the broadcast messages serve the same files to the new users as they did to the old, the same total number of files can be decoded. This can go on as long as every step introduces new $tU_i$ users, which is true for all but the last $(s_i t - d_i)$ steps. Indeed, the term $\hat{H}_{K-(s_i t - d_i)+1}(s_i t, s_i b)$ re-introduces users that previously appeared in $\hat{H}_1(s_i t, s_i b)$. Thus, for every level $i$, only a fraction $\frac{K-(s_i t - d_i)}{K}$ of the cut-set bounds can decode all $p_i$ files using Fano's inequality. We bound this fraction by:

$$\frac{K - (s_i t - d_i)}{K} \geq \begin{cases} \frac{1}{2} & \text{if } s_i t > d_i; \\ 1 & \text{if } s_i t = d_i, \end{cases}$$

and we define $\lambda_i$ to be the right-hand side of the inequality.

By substituing $p_i$ for its value in (16), we get the final bounds:

$$\begin{aligned} bRF + tMF &\geq \sum_{i=1}^{L} \frac{\lambda_i}{s_i} \cdot \min\{(s_i t - d_i + 1)U_i \cdot s_i b, N_i\} \cdot F \\ R &\geq \sum_{i=1}^{L} \lambda_i \cdot \min\left\{(s_i t - d_i + 1)U_i, \frac{N_i}{s_i b}\right\} - \frac{t}{b} M, \end{aligned}$$

which concludes the proof of Lemma 5.

### *B. Proof of approximate optimality (Theorem 4)*

In this section, we will use the information-theoretic lower bounds determined in Lemma 5 to give an upper bound on the ratio between the rate achieved by the memory-sharing scheme and the optimal rate (the "gap"). In order to do that, we must consider a few different cases, in each of which different values are chosen for the parameters $b$, $t$, and $\{s_i\}_i$ defined in Lemma 5. When these parameters are chosen, we will need to evaluate the minimization seen in the expression of the lower bound, for every level. For simplicity, we define, for each level $i$:

$$A_i = \min\left\{(s_i t - d_i + 1)U_i\,,\, \frac{N_i}{s_i b}\right\}. \tag{17}$$

To find or bound the value of $A_i$, we must evaluate the following comparision:

$$bs_i(s_i t - d_i + 1) \overset{?}{\lessgtr} \frac{N_i}{U_i}. \tag{18}$$

Before the main analysis, we consider the case where the number of caches is bounded. Specifically, we consider $K < k_0 = D/\beta$, and we call this "Case 0". Then, we consider the more interesting case where $K$ is unbounded, and divide that into two main regimes. The first, "Case 1", is when the set $I_1$, defined in Definition 1 and (8), is empty; the second, "Case 2", is when it is not empty.

For convenience, we will assume, without loss of generality, that the levels are numbered from most popular to least popular. In other words:

$$U_1/N_1 \geq \cdots \geq U_L/N_L. \tag{19}$$

*1) Case 0: $K < k_0 = D/\beta$:* Recall the refined $M$-feasible partition described in Definition 1 and (8). By (7), we know that the memory given to a level $i_0 \in I_0$ is at most:

$$\alpha_{i_0} M < (2/K)\sqrt{N_{i_0}/U_{i_0}}.$$

However, if $K < D/\beta$, then, for any level $i \in I'$:

$$\begin{aligned}
\alpha_i M &\leq \left(\frac{1}{K} + \frac{\beta}{d_i}\right)\sqrt{\frac{N_i}{U_i}} \\
&= \frac{1}{K}\left(1 + \frac{\beta K}{d_i}\right)\sqrt{\frac{N_i}{U_i}} \\
&< \frac{1}{K}\left(1 + \frac{D}{d_i}\right)\sqrt{\frac{N_i}{U_i}} \\
&\leq \frac{D+1}{K}\sqrt{\frac{N_i}{U_i}} \\
&= \frac{D+1}{2} \cdot \frac{2}{K}\sqrt{\frac{N_i}{U_i}}.
\end{aligned}$$

This suggests that the memory given to a level in $I'$ will not be much larger than that given to a level in $I_0$. As a result, levels in $I'$ are expected to behave similarly to those in $I_0$: they receive so little memory that their impact on the overall transmission rate is not much different from that of the levels in $H$, which get zero memory. The effective result is that the subset $I_1$ "dominates" the set $I$.

With these observations in mind, we next describe a near-equivalent formulation of the achievability scheme, which is more suitable for the case $K < D/\beta$. This formulation will emphasize the fact that $I_1$ is the dominant subset of $I$, by essentially reducing $I$ to $I_1$ and relegating $I_0$ and $I'$ to $H$.

A more precise explanation follows. Find the (unique) level $i^*$ such that:

$$\sum_{i=1}^{i^*-1} \frac{N_i}{d_i} \leq M \leq \sum_{i=1}^{i^*} \frac{N_i}{d_i}.$$

Recall that the levels are numbered from most popular to least popular, as seen in (19).

Partition the set of levels into three sets $(H, I, J)$, which will serve the same purpose as the partition described in Definition 1. We set $H = \{i^*+1, \ldots, L\}$, $I = \{i^*\}$, and $J = \{1, \ldots, i^*-1\}$. We then proceed as usual: the levels in $J$ are fully stored in the caches; the levels in $H$ are given no memory; and the single level in $I$, $i^*$, is given the remaining memory $M - T_J$. The conventional scheme is applied to each level and its corresponding memory, resulting in the following transmission rate:

$$\begin{aligned}
R &= \sum_{h \in H} KU_h + KU_{i^*}\left(1 - \frac{M - T_J}{N_{i^*}/d_{i^*}}\right) \\
&\leq k_0 \left[\sum_{h \in H} U_h + U_{i^*}\left(1 - \frac{M - T_J}{N_{i^*}/d_{i^*}}\right)\right].
\end{aligned} \tag{20}$$

For the lower bounds, consider Lemma 5, with $t = 1$, $s_i = d_i$ for all $i$, and $b = \lceil N_{i^*}/d_{i^*}U_{i^*} \rceil$. Thus $\lambda_i = 1$ for all $i$. We will first analyze the comparisons in (18) for every level.

For $i^*$, we have:

$$bs_{i^*}(s_{i^*}t - d_{i^*} + 1) = bd_{i^*} \geq \frac{N_{i^*}}{U_{i^*}}.$$

For $j \in J$, we have:

$$bs_j(s_j t - d_j + 1) = bd_j \geq \frac{N_{i^*}}{U_{i^*}} \cdot \frac{d_j}{d_{i^*}} \geq \frac{N_{i^*}}{U_{i^*} D} \geq \frac{N_j}{U_j},$$

by regularity condition (3).

Finally, for $h \in H$, we have:

$$bs_h(s_h - d_h + 1) = bd_h \leq 2\frac{N_{i^*}}{U_{i^*}} \cdot \frac{d_h}{d_{i^*}} \leq \frac{N_h}{U_h},$$

again by the regularity condition (3).

Putting these together, we determine the value of $A_i$, defined in (17), for each $i$ and get the following lower bound on the optimal rate:

$$\begin{aligned} R^*(M) &\geq \sum_{h \in H} U_h + \frac{N_{i^*}}{d_{i^*} b} + \sum_{j \in J} \frac{N_j}{d_j b} - \frac{M}{b} \\ &= \sum_{h \in H} U_h + \frac{N_{i^*}/d_{i^*} - (M - T_J)}{b} \\ &\geq \sum_{h \in H} U_h + \frac{N_{i^*}/d_{i^*} - (M - T_J)}{2N_{i^*}/d_{i^*} U_{i^*}} \\ &\geq \frac{1}{2}\left[\sum_{h \in H} U_h + U_{i^*}\left(1 - \frac{M - T_J}{N_{i^*}/d_{i^*}}\right)\right]. \end{aligned} \tag{21}$$

By combining (20) with (21), we get the following bound on the gap:

$$\frac{R(M)}{R^*(M)} \leq 2k_0 = 2D/\beta = 396D. \tag{22}$$

*2) Case 1: $I_1 = \emptyset$:* When $I_1 = \emptyset$, the rate achieved by the memory-sharing scheme is, according to Lemma 3:

$$R(M) \leq \sum_{h \in H} KU_h + \frac{2S_I^2}{M - T_J + V_I}. \tag{23}$$

For technical reasons, the lower bounds analysis for this case has to be broken down into two subcases, depending on whether or not the set $J$ is empty.

Since we will be dealing with many floors and ceilings, here are a few remarks on these operations.

- If $n \geq 1$ is an integer, then $x \geq n \implies \lfloor x \rfloor \geq n$;
- For all $x > 0$, $\lfloor x \rfloor \geq x - 1$ and $\lceil x \rceil \leq x + 1$;
- If $x \geq 1$, then $\lfloor x \rfloor \geq x/2$ and $\lceil x \rceil \leq 2x$;
- If $a \leq x \leq b$, then $\lfloor a \rfloor \leq \lfloor x \rfloor \leq \lfloor b \rfloor$ and $\lceil a \rceil \leq \lceil x \rceil \leq \lceil b \rceil$.

*a) Case 1a: $J \neq \emptyset$:* Recall that $\tilde{M} = (M - T_J + V_I)/S_I$. Consider Lemma 5, with the following parameters:

$$\begin{aligned} & & t &= 1; \\ \forall h \in H, & & s_h &= \left\lfloor \frac{1}{8} K \right\rfloor; \\ \forall i \in I, & & s_i &= \left\lfloor \frac{\sqrt{N_i/U_i}}{8\tilde{M}} \right\rfloor; \\ \forall j \in J, & & s_j &= d_j; \\ & & b &= \left\lfloor \delta \tilde{M}^2 \right\rfloor, \end{aligned}$$

where $\delta = D/\beta$. Notice that $\lambda_j = 1$ for all $j \in J$, and $\lambda_i \geq \frac{1}{2}$ for all other levels $i$.

The first step is to verify that these parameters satisfy all of the constraints imposed on them in Lemma 5. For ease of reference, we repeat the constraints here:

- $t \in \{1, \ldots, K\}$;
- $s_i \in \mathbb{N}^+$ such that $s_i t \in \{d_i, \ldots, \lfloor K/2 \rfloor\}$, for any level $i$;
- $b \in \mathbb{N}^+$.

The parameters $t$ and $s_j$, $j \in J$, can be easily seen to satisfy theirs. Also, it follows from Definition 1 and from regularity condition (1) that, for some $j \in J$:

$$\begin{aligned}\delta\tilde{M}^2 &\geq \delta \cdot \left(\frac{1}{d_j} + \frac{1}{K}\right)^2 \frac{N_j}{U_j} \geq D/\beta \cdot \left(\frac{1}{d_j} + \frac{1}{K}\right)^2 \cdot K \\ &\geq k_0 \cdot \left(\frac{1}{D}\right)^2 \cdot k_0 = \left(\frac{k_0}{D}\right)^2 \geq 1,\end{aligned}$$

and therefore $b \geq 1$.

As for $s_h$, $h \in H$: we have $s_h \leq \lfloor K/2 \rfloor$ trivially, and $s_h \geq d_h$ is true because $K \geq k_0 = D/\beta \geq 8d_h$.

Finally, consider $i \in I$. First, recall that $I_1 = \emptyset$, and therefore $i \notin I_1$. Then, using Definition 1, we have:

$$\begin{aligned}\frac{\sqrt{N_i/U_i}}{8\tilde{M}} &\geq \frac{1}{8} \cdot \frac{1}{\frac{\beta}{d_i} + \frac{1}{K}} = \frac{1}{8(\beta + d_i/K)} \cdot d_i \\ &\geq \frac{1}{16\beta} \cdot d_i \geq d_i,\end{aligned}$$

which implies $s_i \geq d_i$. Also, we have:

$$s_i = \left\lfloor \frac{\sqrt{N_i/U_i}}{8\tilde{M}} \right\rfloor \leq \left\lfloor \frac{1}{8} \cdot \frac{1}{\frac{1}{K}} \right\rfloor \leq \lfloor K/2 \rfloor .$$

Thus all the parameters satisfy their constraints.

Next, we will compute bounds on the $A_i$ terms by evaluating the comparison (18) for all levels $i$.

For $h \in H$, the comparison (18) gives:

$$bs_h(s_h t - dh + 1) \leq bs_h^2 t \leq \delta\tilde{M}^2 \cdot \frac{1}{64}K^2 \leq \frac{\delta}{64} \cdot \frac{N_h}{U_h}.$$

Thus:

$$\begin{aligned}A_h &\geq \min\left\{1, \frac{64}{\delta}\right\} \cdot (s_h t - d_h + 1)U_h \\ &= \min\left\{1, \frac{64\beta}{D}\right\} \cdot (s_h t - d_h + 1)U_h \\ &\geq \frac{64\beta}{D} \cdot \left(\frac{1}{8}K - d_h\right) U_h \\ &\geq \frac{64\beta}{D} \cdot \left(\frac{1}{8} - \beta\right) \cdot KU_h \\ &\geq \frac{8\beta}{D} \cdot (1 - 8\beta) \cdot KU_h.\end{aligned}$$

Consider now $i \in I$. The comparison (18) gives:

$$bs_i(s_i t - d_i + 1) \leq bs_i^2 t \leq \delta\tilde{M}^2 \cdot \frac{1}{64}\frac{N_i/U_i}{\tilde{M}^2} = \frac{\delta}{64} \cdot \frac{N_i}{U_i}.$$

Therefore,

$$\begin{aligned}A_i &\geq \min\left\{1, \frac{64}{\delta}\right\} \cdot (s_i t - d_i + 1)U_i \\ &= \frac{64\beta}{D} \cdot (s_i t - d_i + 1)U_i \\ &\geq \frac{64\beta}{D} \cdot \left(\frac{\sqrt{N_i/U_i}}{8\tilde{M}} - d_i\right) U_i \\ &= \frac{64\beta}{D} \cdot \left(\frac{1}{8} - d_i \cdot \frac{\tilde{M}}{\sqrt{N_i/U_i}}\right) \cdot \frac{\sqrt{N_i U_i}}{\tilde{M}} \\ &\geq \frac{64\beta}{D} \cdot \left(\frac{1}{8} - \left(\beta + \frac{d_i}{K}\right)\right) \cdot \frac{\sqrt{N_i U_i}}{\tilde{M}} \\ &\geq \frac{64\beta}{D} \cdot \left(\frac{1}{8} - 2\beta\right) \cdot \frac{\sqrt{N_i U_i}}{\tilde{M}} \\ &\geq \frac{8\beta}{D} \cdot (1 - 16\beta) \cdot \frac{\sqrt{N_i U_i}}{\tilde{M}}.\end{aligned}$$

Finally, let us look at $j \in J$. The comparison in (18) gives:

$$\begin{aligned} bs_j(s_j t - d_j + 1) &= bd_j \\ &\geq \frac{1}{2}\delta \tilde{M}^2 \cdot d_j \\ &\geq \frac{1}{2}(D/\beta)\left(\frac{1}{d_j} + \frac{1}{K}\right)^2 \frac{N_j}{U_j} d_j \\ &\geq \frac{1}{2\beta} d_j^2 \cdot \left(\frac{1}{d_j}\right)^2 \cdot \frac{N_j}{U_j} \\ &\geq \frac{N_j}{U_j}. \end{aligned}$$

Therefore:

$$A_j = \frac{N_j}{d_j b}.$$

We can now use the values of $\{A_i\}_i$ to lower-bound the optimal rate as in Lemma 5. Recall also the values of $\lambda_i$ mentioned at the beginning of this section.

$$\begin{aligned} R^*(M) &\geq \sum_{h \in H} \frac{1}{2} \cdot \frac{8\beta}{D} \cdot (1 - 8\beta) \cdot KU_h \\ &\quad + \sum_{i \in I} \frac{1}{2} \cdot \frac{8\beta}{D} \cdot (1 - 16\beta) \cdot \frac{\sqrt{N_i U_i}}{\tilde{M}} \\ &\quad + \sum_{j \in J} \frac{N_j}{d_j b} - \frac{M}{b} \\ &= \sum_{h \in H} \frac{4\beta}{D} \cdot (1 - 8\beta) \cdot KU_h \\ &\quad + \frac{4\beta}{D} \cdot (1 - 16\beta) \cdot \frac{S_I \cdot \sum_{i \in I} \sqrt{N_i U_i}}{M - T_J + V_I} \\ &\quad - \frac{M - T_J}{b} \\ &\geq \sum_{h \in H} \frac{4\beta}{D} \cdot (1 - 8\beta) \cdot KU_h \\ &\quad + \frac{4\beta}{D} \cdot (1 - 16\beta) \cdot \frac{S_I^2}{M - T_J + V_I} \\ &\quad - \frac{M - T_J + V_I}{\frac{1}{2}\delta (M - T_J + V_I)^2 / S_I^2} \\ &= \sum_{h \in H} \frac{4\beta}{D} \cdot (1 - 8\beta) \cdot KU_h \\ &\quad + \left[\frac{4\beta}{D} \cdot (1 - 16\beta) - \frac{2}{\delta}\right] \cdot \frac{S_I^2}{M - T_J + V_I} \\ &= \sum_{h \in H} \frac{\beta}{D} \cdot (4 - 32\beta) \cdot KU_h \\ &\quad + \frac{\beta}{D}(1 - 32\beta) \cdot \frac{2S_I^2}{M - T_J + V_I}. \end{aligned} \tag{24}$$

Combining (24) with (23), we get:

$$\frac{R(M)}{R^*(M)} \leq \frac{D/\beta}{1 - 32\beta} \leq 237D. \tag{25}$$

*b) Case 1b: $J = \emptyset$:* We know that $I$ is never an empty set. Since $J = \emptyset$, then the most popular level is in $I$. By (19), that level is level 1, and hence $1 \in I$. Moreover, we note that $T_J = 0$ when $J = \emptyset$, and thus $\tilde{M} = (M + V_I)/S_I$.

Consider again Lemma 5, and define the following parameters:

$$\begin{aligned} t &= \left\lfloor \frac{1}{32} \frac{\sqrt{N_1/U_1}}{\tilde{M}} \right\rfloor; \\ \forall h \in H, \qquad s_h &= \left\lfloor 2 \frac{K\tilde{M}}{\sqrt{N_1/U_1}} \right\rfloor; \\ \forall i \ \in I, \qquad s_i &= \left\lfloor 2 \sqrt{\frac{N_i/U_i}{N_1/U_1}} \right\rfloor; \\ b &= \left\lfloor 8\tilde{M}\sqrt{N_1/U_1} \right\rfloor. \end{aligned}$$

Thus, we can only say that $\lambda_i \geq \frac{1}{2}$ for all levels $i$.

As before, we must first verify that these parameters satisfy all their constraints. Let us start with $b$:

$$8\tilde{M}\sqrt{N_1/U_1} \geq 8 \cdot \frac{1}{K}\sqrt{N_1/U_1} \cdot \sqrt{N_1/U_1} = \frac{8N_1}{KU_1} \geq 1,$$

by regularity condition (1), and therefore $b \geq 1$.

Next, we will examine $t$ and the $\{s_i\}_i$. It suffices to show $t \geq 1$, $s_i \geq 1$, $s_i t \geq d_i$, and $s_i t \leq \lfloor K/2 \rfloor$ for all $i$. We can see that $t \geq 1$ because:

$$\frac{\sqrt{N_1/U_1}}{32\tilde{M}} \geq \frac{1/32}{\frac{\beta}{d_1} + \frac{1}{K}} \geq \frac{1}{64\beta} \cdot d_1 \geq 1,$$

where the first inequality follows from $1 \in I \setminus I_1$ (recall $I_1 = \emptyset$) and Definition 1 combined with (8).

For $h \in H$, we have $2K\tilde{M}/\sqrt{N_1/U_1} \geq 1$ directly from Definition 1, and thus $s_h \geq 1$. Moreover,

$$s_h t \geq \frac{1}{64} \cdot \frac{\sqrt{N_1/U_1}}{\tilde{M}} \cdot \frac{K\tilde{M}}{\sqrt{N_1/U_1}} = \frac{K}{64} \geq D \geq d_h,$$

and,

$$s_h t \leq \frac{1}{32} \cdot \frac{\sqrt{N_1/U_1}}{\tilde{M}} \cdot 2\frac{K\tilde{M}}{\sqrt{N_1/U_1}} = K/16 \leq \lfloor K/2 \rfloor.$$

Consider now $i \in I$. Because level 1 is the most popular level, then $N_i/U_i \geq N_1/U_1$ for all $i$, and hence $s_i \geq 1$. Moreover:

$$s_i t \geq \frac{\sqrt{N_i/U_i}}{64\tilde{M}} \geq \frac{1/64}{\frac{\beta}{d_i} + \frac{1}{K}} \geq \frac{1}{128\beta} d_i \geq d_i.$$

Finally,

$$s_i t \leq \frac{\sqrt{N_i/U_i}}{16\tilde{M}} \leq K/16 \leq \lfloor K/2 \rfloor.$$

Thus, all the parameters satisfy their conditions.

We now move to the comparisons in (18) which allow us to give bounds on $A_i$ for all $i$.

Consider $h \in H$. Note that:

$$\begin{aligned} b s_h (s_h t - d_h + 1) &\leq b s_h^2 t \\ &\leq 8\tilde{M}\sqrt{N_1/U_1} \cdot \left[ \frac{2K\tilde{M}}{\sqrt{N_1/U_1}} \right]^2 \cdot \frac{\sqrt{N_1/U_1}}{32\tilde{M}} \\ &= \left[ K\tilde{M} \right]^2 \\ &\leq \frac{N_h}{U_h}. \end{aligned}$$

Therefore,

$$\begin{aligned}
A_h &= (s_h t - d_h + 1)U_h \\
&\geq \left[\left(\frac{K\tilde{M}}{\sqrt{N_1/U_1}}\right)\left(\frac{\sqrt{N_1/U_1}}{32\tilde{M}} - 1\right) - (d_h - 1)\right] U_h \\
&\geq \left[\left(\frac{1}{32} - \frac{\tilde{M}}{\sqrt{N_1/U_1}}\right) - \frac{D-1}{k_0}\right] KU_h \\
&\overset{(a)}{\geq} \left[\frac{1}{32} - \left(\frac{\beta}{d_1} + \frac{1}{K}\right) - \frac{D}{k_0} - \frac{1}{k_0}\right] KU_h \\
&\overset{(b)}{\geq} \left(\frac{1}{32} - 2\beta\right) \cdot KU_h.
\end{aligned}$$

Here, $(a)$ uses Definition 1 with (8) and $1 \in I \setminus I_1$, and $(b)$ uses $d_1 \geq 1$ and $D/k_0 = \beta$.

Now consider $i \in I$. We have:

$$\begin{aligned}
bs_i(s_i t - d_i + 1) &\leq bs_i^2 t \\
&\leq 8\tilde{M}\sqrt{N_1/U_1} \cdot 4\frac{N_i/U_i}{N_1/U_1} \cdot \frac{\sqrt{N_1/U_1}}{32\tilde{M}} \\
&= \frac{N_i}{U_i}
\end{aligned}$$

Therefore,

$$\begin{aligned}
A_i &= (s_i t - d_i + 1)U_i \\
&\geq \left[\sqrt{\frac{N_i/U_i}{N_1/U_1}} \cdot \left(\frac{\sqrt{N_1/U_1}}{32\tilde{M}} - 1\right) - (d_i - 1)\right] U_i \\
&= \left[\frac{1}{32} - \frac{\tilde{M}}{\sqrt{N_1/U_1}} - (d_i - 1) \cdot \frac{\tilde{M}}{\sqrt{N_i/U_i}}\right] \frac{\sqrt{N_i U_i}}{\tilde{M}} \\
&\geq \left[\frac{1}{32} - \left(\frac{\beta}{d_1} + \frac{1}{K}\right) - (d_i - 1) \cdot \left(\frac{\beta}{d_i} + \frac{1}{K}\right)\right] \frac{\sqrt{N_i U_i}}{\tilde{M}} \\
&= \left[\frac{1}{32} - 2\beta - \left(\frac{1}{d_1} - \frac{1}{d_i}\right)\beta\right] \frac{\sqrt{N_i U_i}}{\tilde{M}} \\
&\geq \left(\frac{1}{32} - 3\beta\right) \frac{\sqrt{N_i U_i}}{\tilde{M}}.
\end{aligned}$$

Finally, we will also give an upper bound to the following useful quantity:

$$\begin{aligned}
\frac{tM}{b} &\leq \frac{\sqrt{N_1/U_1}}{32\tilde{M}} \cdot \frac{2}{8\tilde{M}\sqrt{N_1/U_1}} \cdot M \\
&\leq \frac{S_I^2}{128(M + V_I)^2} \cdot (M + V_I) \\
&= \frac{S_I^2}{128(M + V_I)}.
\end{aligned}$$

Combining the $A_i$ terms, along with the $\lambda_i$ factors, we get the following lower bound:

$$\begin{aligned}
R^*(M) &\geq \sum_{h\in H} \frac{1}{2}\cdot\left(\frac{1}{32}-2\beta\right)\cdot KU_h \\
&\quad + \sum_{i\in I} \frac{1}{2}\cdot\left(\frac{1}{32}-3\beta\right)\cdot\frac{\sqrt{N_iU_i}}{\tilde{M}} \\
&\quad - \frac{tM}{b} \\
&\geq \sum_{h\in H} \frac{1}{2}\cdot\left(\frac{1}{32}-2\beta\right)\cdot KU_h \\
&\quad + \frac{1}{2}\cdot\left(\frac{1}{32}-3\beta\right)\cdot\frac{S_I^2}{M+V_I} \\
&\quad - \frac{S_I^2}{128(M+V_I)} \\
&= \sum_{h\in H} \frac{1}{2}\cdot\left(\frac{1}{32}-2\beta\right)\cdot KU_h \\
&\quad + \left[\frac{1}{2}\cdot\left(\frac{1}{32}-3\beta\right)-\frac{1}{128}\right]\cdot\frac{S_I^2}{M+V_I}. \\
&= \sum_{h\in H} \frac{1}{64}(1-64\beta)\cdot KU_h \\
&\quad + \frac{1}{256}(1-192\beta)\cdot\frac{2S_I^2}{M+V_I}. \qquad (26)
\end{aligned}$$

Combining (26) with (23):

$$\frac{R(M)}{R^*(M)} \leq \frac{256}{1-192\beta} = 8448. \qquad (27)$$

*3) Case 2: $I_1 \neq \emptyset$:* In this section, there exists at least one level in $I_1$. However, by Proposition 2, the set $I_1$ cannot contain more than one level. Therefore, in this section we have $I_1 = \{i_1\}$ for some level $i_1$.

Using Lemma 3, we can write the achievable rate as follows:

$$\begin{aligned}
R(M) &\leq \sum_{h\in H} KU_h + \frac{2S_IS_{I\setminus I_1}}{M-T_J+V_I} \\
&\quad + \frac{1}{\beta}d_{i_1}U_{i_1}\left(1-\frac{M-T_J}{N_{i_1}/d_{i_1}}\right) \\
&\quad + \frac{1}{\beta}d_{i_1}U_{i_1}\cdot\frac{S_{I\setminus I_1}}{\sqrt{N_{i_1}U_{i_1}}}. \qquad (28)
\end{aligned}$$

Let $\gamma_1 = 2.965$ and $\gamma_2 = 0.482$. Consider Lemma 5 with the following parameters.

$$\begin{aligned}
& & t &= 1; \\
\forall h \in H, & & s_h &= \lfloor 2\beta K \rfloor; \\
\forall i \in I\setminus I_1, & & s_i &= \left\lfloor \gamma_1\beta\frac{\sqrt{N_i/U_i}}{\tilde{M}} + \gamma_2 d_{i_1}\sqrt{\frac{U_{i_1}/N_{i_1}}{U_i/N_i}} \right\rfloor; \\
& & s_{i_1} &= d_{i_1}; \\
\forall j \in J, & & s_j &= d_j; \\
& & b &= \left\lceil \frac{N_{i_1}}{d_{i_1}U_{i_1}} \right\rceil.
\end{aligned}$$

Thus we have $\lambda_{i_1} = \lambda_j = 1$ for all $j \in J$, and $\lambda_i, \lambda_h \geq \frac{1}{2}$ for $i \in I$ and $h \in H$.

Recall the following inequalities from Definition 1 regarding $i \in I$:

$$\frac{1}{K}\sqrt{\frac{N_i}{U_i}} \leq \tilde{M} \leq \left(\frac{1}{K}+\frac{1}{d_i}\right)\sqrt{\frac{N_i}{U_i}}.$$

Thus, for any $i, j \in I$, we have:

$$\sqrt{\frac{U_i/N_i}{U_j/N_j}} \leq \frac{K}{d_i} + 1.$$

Moreover, for $h \in H$ and $i_1 \in I_1$:

$$\left(\frac{1}{K} + \frac{\beta}{d_{i_1}}\right) \leq \tilde{M} \leq \frac{1}{K}\sqrt{\frac{N_h}{U_h}},$$

and hence:

$$\sqrt{\frac{U_{i_1}/N_{i_1}}{U_h/N_h}} \geq \beta\frac{K}{d_{i_1}} + 1.$$

We will now verify that the parameters satisfy their constraints. The parameters $t$, $s_{i_1}$, and $s_j$ trivially satisfy all their constraints. Moreover, $N_{i_1} \geq KU_{i_1} \geq d_{i_1}U_{i_1}$, which implies $b \geq 1$. Regarding $s_h$:

$$K/2 \geq 2\beta K \geq 2\beta \cdot D/\beta = 2D \geq d_h,$$

and hence $d_h \leq s_h \leq \lfloor K/2 \rfloor$. We are hence only left with $s_i$.

Consider the following:

$$\begin{aligned}
&\gamma_1\beta\frac{\sqrt{N_i/U_i}}{\tilde{M}} + \gamma_2 d_{i_1}\sqrt{\frac{U_{i_1}/N_{i_1}}{U_i/N_i}} \\
&\geq \gamma_1\beta\frac{\sqrt{N_i/U_i}}{\tilde{M}} \geq \frac{\gamma_1\beta}{\beta/d_i + 1/K} \\
&\geq \frac{\gamma_1\beta}{2\beta}d_i \geq d_i,
\end{aligned}$$

and hence $s_i \geq d_i$. Furthermore,

$$\begin{aligned}
&\gamma_1\beta\frac{\sqrt{N_i/U_i}}{\tilde{M}} + \gamma_2 d_{i_1}\sqrt{\frac{U_{i_1}/N_{i_1}}{U_i/N_i}} \\
&\leq \gamma_1\beta K + \gamma_2 d_{i_1}\left(\frac{K}{d_{i_1}} + 1\right) \\
&= \left[\gamma_1\beta + \gamma_2\left(1 + \frac{d_{i_1}}{K}\right)\right]K \\
&\leq \left[\gamma_2 + (\gamma_1 + \gamma_2)\beta\right]K \\
&\leq K/2,
\end{aligned}$$

and hence $s_i = s_i t \leq \lfloor K/2 \rfloor$.

We will now evaluate the comparisons in (18) to give bounds on the $A_i$ terms, for every $i$.

Let us start with $h \in H$. We have $bs_h(s_h t - d_h + 1) \leq bs_h^2 t = bs_h^2$, and:

$$\begin{aligned}
bs_h^2 &\leq \left(\frac{N_{i_1}}{d_{i_1}U_{i_1}} + 1\right) \cdot 4\beta^2K^2 \\
&\leq \frac{N_{i_1}}{d_{i_1}U_{i_1}}\left(1 + \frac{d_{i_1}}{K}\right) \cdot 4\beta^2K^2 \\
&\overset{(a)}{\leq} \frac{N_h}{d_{i_1}U_h}\frac{1 + d_{i_1}/K}{(\beta + d_{i_1}/K)^2} \cdot \frac{d_{i_1}^2}{K^2} \cdot 4\beta^2K^2 \\
&\overset{(b)}{\leq} \frac{N_h}{U_h} \cdot \frac{1}{\beta^2} \cdot d_{i_1} \cdot 4\beta^2 \\
&\leq 4D \cdot \frac{N_h}{U_h}.
\end{aligned}$$

The inequality labeled $(a)$ comes from Definition 1 and (8), and:

$$\left(\frac{\beta}{d_{i_1}} + \frac{1}{K}\right)\sqrt{\frac{N_{i_1}}{U_{i_1}}} < M \leq \frac{1}{K}\sqrt{\frac{N_h}{U_h}}.$$

As for the inequality $(b)$, it is due to:

$$\forall x \geq 0, \quad \frac{1 + x}{(\beta + x)^2} = \frac{1}{\beta + x} \cdot \frac{1 + x}{\beta + x} \leq \frac{1}{\beta} \cdot \frac{1}{\beta},$$

because $x \to \frac{1+x}{\beta+x}$ is decreasing for $x \geq 0$. Hence,

$$\begin{aligned} A_h &\geq \min\left\{1, \frac{1}{4D}\right\} \cdot (s_h t - d_h + 1) U_h \\ &\geq \frac{1}{4D} \cdot (2\beta K - d_h)\, U_h \\ &\geq \frac{1}{4D} \cdot (2\beta - \beta)\, K U_h \\ &= \frac{\beta}{4D} \cdot K U_h. \end{aligned}$$

Now we consider $i \in I \setminus I_1$. We have:

$$\begin{aligned} bs_i^2 &\leq \frac{N_{i_1}}{d_{i_1} U_{i_1}} \left(1 + \frac{d_{i_1}}{K}\right) \left(\gamma_1 \beta \frac{\sqrt{N_i/U_i}}{\tilde{M}} + \gamma_2 d_{i_1} \sqrt{\frac{U_{i_1}/N_{i_1}}{U_i/N_i}}\right)^2 \\ &\leq \frac{N_{i_1}}{d_{i_1} U_{i_1}} (1+\beta) \left(\gamma_1 beta \frac{\sqrt{N_{i_1}/U_{i_1}}}{\tilde{M}} + \gamma_2 d_{i_1}\right)^2 \cdot \frac{U_{i_1}/N_{i_1}}{U_i/N_i} \\ &\leq \frac{N_i}{d_{i_1} U_i} (1+\beta) \left(\frac{\gamma_1 \beta}{\beta/d_{i_1} + 1/K} + \gamma_2 d_{i_1}\right)^2 \\ &\leq \frac{N_i}{d_{i_1} U_i} (1+\beta) \left(\frac{\gamma_1 \beta}{\beta} d_{i_1} + \gamma_2 d_{i_1}\right)^2 \\ &= \frac{N_i}{d_{i_1} U_i} (1+\beta) d_{i_1}^2 \left(\gamma_1 + \gamma_2\right)^2 \\ &\leq \frac{N_i}{U_i} \cdot 12(1+\beta) D. \end{aligned}$$

Therefore, if we define $c = \frac{1}{12(1+\beta)D}$, then:

$$\begin{aligned} A_i &\geq c \cdot (s_i t - d_i + 1) U_i \\ &\geq c \cdot \left(\gamma_1 \beta \frac{\sqrt{N_i/U_i}}{\tilde{M}} + \gamma_2 d_{i_1} \sqrt{\frac{U_{i_1}/N_{i_1}}{U_i/N_i}} - d_i\right) U_i \\ &= c \left[\left(\gamma_1 \beta \frac{\sqrt{N_i/U_i}}{\tilde{M}} - d_i\right) U_i + \gamma_2 d_{i_1} \sqrt{\frac{U_{i_1}/N_{i_1}}{U_i/N_i}} \cdot U_i\right] \\ &= c \left[\left(\gamma_1 \beta - d_i \frac{\tilde{M}}{\sqrt{N_i/U_i}}\right) \frac{\sqrt{N_i U_i}}{\tilde{M}} + \gamma_2 d_{i_1} U_{i_1} \frac{\sqrt{N_i U_i}}{\sqrt{N_{i_1} U_{i_1}}}\right] \\ &\geq c \left[(\gamma_1 \beta - 2\beta) \frac{\sqrt{N_i U_i}}{\tilde{M}} + \gamma_2 d_{i_1} U_{i_1} \frac{\sqrt{N_i U_i}}{\sqrt{N_{i_1} U_{i_1}}}\right] \\ &= c \left[(\gamma_1 - 2)\beta \frac{\sqrt{N_i U_i}}{\tilde{M}} + \gamma_2 d_{i_1} U_{i_1} \frac{\sqrt{N_i U_i}}{\sqrt{N_{i_1} U_{i_1}}}\right] \\ &\geq c \cdot \min\left\{\frac{1}{2}\gamma_1 - 1, \gamma_2\right\} \cdot \beta \\ &\quad \cdot \left[\frac{2\sqrt{N_i U_i}}{\tilde{M}} + \frac{1}{\beta} d_{i_1} U_{i_1} \frac{\sqrt{N_i U_i}}{\sqrt{N_{i_1} U_{i_1}}}\right] \\ &= 0.482 c\beta \cdot \left[\frac{2\sqrt{N_i U_i}}{\tilde{M}} + \frac{1}{\beta} d_{i_1} U_{i_1} \frac{\sqrt{N_i U_i}}{\sqrt{N_{i_1} U_{i_1}}}\right]. \end{aligned}$$

Finally, we have:

$$bs_{i_1}(s_{i_1} t - d_{i_1} + 1) = bd_{i_1} \geq \frac{N_{i_1}}{U_{i_1}},$$

and, likewise for $j \in J$:

$$bs_j(s_j t - d_j + 1) = bd_j \geq \frac{d_j}{d_{i_1}} \frac{N_{i_1}}{U_{i_1}} \geq \frac{N_j}{U_j},$$

because of the regularity condition (3). Thus:

$$A_{i_1} = \frac{N_{i_1}/d_{i_1}}{b}; \quad A_j = \frac{N_j/d_j}{b}.$$

Combining the $A_i$ terms, and keeping in mind the $\lambda_i$ terms, we get:

$$\begin{aligned}
R^*(M) \geq & \sum_{h\in H} \frac{1}{2}\cdot\frac{\beta}{4D}\cdot KU_h \\
& + \sum_{i\in I\setminus I_1} \frac{1}{2}\cdot\frac{0.482\beta}{12(1+\beta)D}\cdot\left[\frac{2\sqrt{N_iU_i}}{\tilde{M}} + \frac{1}{\beta}d_{i_1}U_{i_1}\frac{\sqrt{N_iU_i}}{N_{i_1}U_{i_1}}\right] \\
& + \frac{N_{i_1}/d_{i_1} + \sum_{j\in J} N_j/d_j - M}{b} \\
\geq & \sum_{h\in H} \frac{\beta}{8D}\cdot KU_h \\
& + \frac{0.241\beta}{12(1+\beta)D}\left[\frac{2S_IS_{I\setminus I_1}}{M - T_J + V_I} + \frac{1}{\beta}d_{i_1}U_{i_1}\frac{S_{I\setminus I_1}}{N_{i_1}U_{i_1}}\right] \\
& + \frac{N_{i_1}/d_{i_1} + \sum_{j\in J} N_j/d_j - M}{2N_{i_1}/(d_{i_1}U_{i_1})} \\
= & \sum_{h\in H} \frac{\beta}{8D}\cdot KU_h \\
& + \frac{0.241\beta}{12(1+\beta)D}\left[\frac{2S_IS_{I\setminus I_1}}{M - T_J + V_I} + \frac{1}{\beta}d_{i_1}U_{i_1}\frac{S_{I\setminus I_1}}{N_{i_1}U_{i_1}}\right] \\
& + \frac{1}{2}U_{i_1}\left(1 - \frac{M - T_J}{N_{i_1}/d_{i_1}}\right).
\end{aligned} \tag{29}$$

Combining (29) with (28), we get:

$$\frac{R(M)}{R^*(M)} \leq \max\left\{\frac{8D}{\beta}, \frac{12(1+\beta)D}{0.241\beta}, \frac{2D}{\beta}\right\} \leq 9909D. \tag{30}$$

*4) Final gap:* By combining the inequalities from all the cases above, i.e., inequalities (22), (25), (27), and (30), we get the following final result:

$$\frac{R(M)}{R^*(M)} \leq 9909\cdot D,$$

which concludes the proof of Theorem 4.

### C. Proofs of some useful results

*Proof of Lemma 3:*

We will prove the lemma for each of the five sets in the refined $M$-feasible partition. First, recall that the achievable rate $R^{\text{SL}}(\cdot)$ for the single-level setup is defined in Theorem 1.

For $h \in H$:

$$R_h(M) = R^{\text{SL}}(0, K, N_h, U_h, d_h) = KU_h.$$

For $i \in I_0$:

$$\begin{aligned}
R_i(M) &= R^{\text{SL}}(\alpha_i M, K, N_i, U_i, d_i) \\
&\leq R^{\text{SL}}(0, K, N_i, U_i, d_i) \\
&= KU_i \\
&\leq \frac{2}{\tilde{M}}\sqrt{\frac{N_i}{U_i}}\cdot U_i \\
&= \frac{2S_I\sqrt{N_iU_i}}{M - T_J + V_I}.
\end{aligned}$$

$$
\begin{aligned}
\Delta &= \left(1-\frac{\alpha_i M}{N_i/d_i}\right)-\left(1-\frac{M-T_J}{N_i/d_i}\right)=\left(1-\frac{\frac{\sqrt{N_iU_i}}{S_I}(M-T_J+V_I)-\frac{N_i}{K}}{N_i/d_i}\right)-\left(1-\frac{M-T_J}{N_i/d_i}\right)\\
&= \frac{M-T_J}{N_i/d_i}-\frac{\sqrt{N_iU_i}(M-T_J)}{S_IN_i/d_i}-\frac{\sqrt{N_iU_i}V_I}{S_IN_i/d_i}+\frac{d_i}{K}=\frac{M-T_J}{N_i/d_i}\left(1-\frac{\sqrt{N_iU_i}}{S_I}\right)-\frac{\sqrt{N_iU_i}V_{I\setminus I_1}}{S_IN_i/d_i}-\frac{\sqrt{N_iU_i}\cdot\frac{N_i}{K}}{S_IN_i/d_i}+\frac{d_i}{K}\\
&= \frac{M-T_J}{N_i/d_i}\cdot\frac{S_{I\setminus I_1}}{S_I}-\frac{d_iV_{I\setminus I_1}}{\sqrt{N_i/U_i}S_I}+\frac{d_i}{K}\left(1-\frac{\sqrt{N_iU_i}}{S_I}\right)=\frac{d_i}{N_i}(M-T_J)\frac{S_{I\setminus I_1}}{S_I}-\frac{d_iV_{I\setminus I_1}}{\sqrt{N_i/U_i}S_I}+\frac{d_i}{K}\cdot\frac{S_{I\setminus I_1}}{S_I}\\
&\overset{(a)}{\leq} \frac{d_i}{N_i}\left[\left(\frac{1}{K}+\frac{1}{d_i}\right)\sqrt{\frac{N_i}{U_i}}S_I-V_I\right]\frac{S_{I\setminus I_1}}{S_I}-\frac{d_iV_{I\setminus I_1}}{\sqrt{N_i/U_i}S_I}+\frac{d_i}{K}\cdot\frac{S_{I\setminus I_1}}{S_I}\\
&= \left(\frac{d_i}{K}+1\right)\frac{S_{I\setminus I_1}}{\sqrt{N_iU_i}}-\frac{d_iV_IS_{I\setminus I_1}}{N_iS_I}-\frac{d_iV_{I\setminus I_1}\sqrt{N_iU_i}}{N_iS_I}+\frac{d_iS_{I\setminus I_1}}{KS_I}\\
&= \left(\frac{d_i}{K}+1\right)\frac{S_{I\setminus I_1}}{\sqrt{N_iU_i}}-\frac{d_i\left(V_{I\setminus I_1}+\frac{N_i}{K}\right)S_{I\setminus I_1}}{N_iS_I}-\frac{d_iV_{I\setminus I_1}\sqrt{N_iU_i}}{N_iS_I}+\frac{d_iS_{I\setminus I_1}}{KS_I}\\
&= \left(\frac{d_i}{K}+1\right)\frac{S_{I\setminus I_1}}{\sqrt{N_iU_i}}-\frac{d_iV_{I\setminus I_1}S_I}{N_iS_I}-\frac{d_iS_{I\setminus I_1}}{KS_I}+\frac{d_iS_{I\setminus I_1}}{KS_I}=\frac{d_i}{K}\cdot\frac{S_{I\setminus I_1}}{\sqrt{N_iU_i}}-\frac{d_iV_{I\setminus I_1}}{N_i}+\frac{S_{I\setminus I_1}}{\sqrt{N_iU_i}}\\
&= \frac{d_i}{K}\sum_{i\in I\setminus I_1}\left(\frac{\sqrt{N_iU_i}}{\sqrt{N_iU_i}}-\frac{N_i}{N_i}\right)+\frac{S_{I\setminus I_1}}{\sqrt{N_iU_i}}=\frac{d_i}{K}\sum_{i\in I\setminus I_1}\frac{\sqrt{N_iU_i}+N_i\sqrt{U_i/N_i}}{\sqrt{N_iU_i}}+\frac{S_{I\setminus I_1}}{\sqrt{N_iU_i}}\\
&= \frac{d_i}{K}\sum_{i\in I\setminus I_1}\frac{N_i\left(\sqrt{U_i/N_i}-\sqrt{N_i/U_i}\right)}{\sqrt{N_iU_i}}+\frac{S_{I\setminus I_1}}{\sqrt{N_iU_i}}\\
&\overset{(b)}{\leq} \frac{S_{I\setminus I_1}}{\sqrt{N_iU_i}}.
\end{aligned}
\tag{31}
$$

---

For $i\in I'$:

$$
\begin{aligned}
R_i(M) &= R^{\mathrm{SL}}(\alpha_iM,K,N_i,U_i,d_i)\\
&\leq U_i\cdot\frac{N_i}{\alpha_iM}\cdot\left(1-\frac{\alpha_iM}{N_i/d_i}\right)\\
&\leq \frac{N_iU_i}{\alpha_iM}\\
&= \frac{N_iU_i}{\sqrt{N_iU_i}\cdot\tilde{M}-N_i/K}\\
&\leq \frac{N_iU_i}{\sqrt{N_iU_i}\cdot\tilde{M}-\frac{1}{2}\sqrt{N_iU_i}\cdot\tilde{M}}\\
&= \frac{2S_I\sqrt{N_iU_i}}{M-T_J+V_I}.
\end{aligned}
$$

For $i\in I_1$:

$$
R_i(M)\leq U_i\cdot\frac{N_i}{\alpha_iM}\cdot\left(1-\frac{\alpha_iM}{N_i/d_i}\right)\leq\frac{1}{\beta}d_iU_i\left(1-\frac{\alpha_iM}{N_i/d_i}\right),
$$

since $i\in I_1\implies\alpha_iM\geq\beta N_i/d_i$ (follows from (8)).

Consider the following difference:

$$
\Delta=\left(1-\frac{\alpha_iM}{N_i/d_i}\right)-\left(1-\frac{M-T_J}{N_i/d_i}\right).
$$

We can show that:

$$
\Delta\leq\frac{S_{I\setminus I_1}}{\sqrt{N_iU_i}}.
$$

The full derivation is given in (31). In the calculations in (31), $(a)$ is due to the definition of $I_1$, while $(b)$ is due to the fact that the (unique) level in $I_1$ is more popular than any other level in $i'\in I$, and thus:

$$
\sqrt{\frac{U_{i'}}{N_{i'}}}\leq\sqrt{\frac{U_i}{N_i}}.
$$

| | $M_1$ | $M_2$ | $M_3$ | $M_4$ | $M$ |
|---|---|---|---|---|---|
| $H$ | $\{1,2\}$ | $\{2\}$ | $\emptyset$ | $\emptyset$ | $\emptyset$ |
| $I$ | $\emptyset$ | $\{1\}$ | $\{1,2\}$ | $\{2\}$ | $\emptyset$ |
| $J$ | $\emptyset$ | $\emptyset$ | $\emptyset$ | $\{1\}$ | $\{1,2\}$ |

| | $M'_1$ | $M'_2$ | $M'_3$ | $M'_4$ | $M$ |
|---|---|---|---|---|---|
| $H$ | $\{1,2\}$ | $\{2\}$ | $\{2\}$ | $\emptyset$ | $\emptyset$ |
| $I$ | $\emptyset$ | $\{1\}$ | $\emptyset$ | $\{2\}$ | $\emptyset$ |
| $J$ | $\emptyset$ | $\emptyset$ | $\{1\}$ | $\{1\}$ | $\{1,2\}$ |

Fig. 16. Two examples of how an $M$-feasible partition $(H, I, J)$ could evolve when there are two levels $\{1, 2\}$. In the second case, it will turn out that $M'_2 = M'_3$. The reasoning behind this is that between $M'_2$ and $M'_3$, all levels have a fixed memory: level 2 has memory 0 while level 1 has memory $N_1/d_1$. Thus an increase in the overall memory between $M'_2$ and $M'_3$ would be wasted. Moreover, $M_1 = M'_1 = 0$ for similar reasons, and $M_4 = M'_4 = N_1/d_1 + N_2/d_2$ because that is the point at which both levels can be completely stored.

Using the bound on $\Delta$, we can rewrite $R_i(M)$ as follows:

$$\begin{aligned} R_i(M) &\le \frac{1}{\beta} d_i U_i \left(1 - \frac{\alpha_i M}{N_i/d_i}\right) \\ &= \frac{1}{\beta} d_i U_i \left(1 - \frac{M - T_J}{N_i/d_i} + \Delta\right) \\ &\le \frac{1}{\beta} d_i U_i \left(1 - \frac{M - T_J}{N_i/d_i}\right) + \frac{1}{\beta} d_i U_i \cdot \frac{S_{I_0} + S_{I'}}{\sqrt{N_i U_i}}. \end{aligned}$$

Finally, for $j \in J$:

$$R_j(M) = R^{\text{SL}}(N_j/d_j, K, N_j, U_j, d_j) = 0.$$

This completes the proof of the lemma. ■

*Proof of Lemma 2:*

We prove the existence of an $M$-feasible partition by construction. In fact, we will use Algorithm 1 presented in Section V-A as the proof. In what follows, we prove the correctness of the algorithm, and then conclude the proof of Lemma 2.

The idea behind the algorithm is as follows. When the memory is 0, all levels are in the set $H$.[10] As the memory is increased, levels will start moving from $H$ to $I$ as they gain more memory, and then finally from $I$ to $J$ as they get completely stored in the caches. Thus, the range of memory values $[0, \infty)$ can be divided into $2L + 1$ intervals; in each interval, the partition $(H, I, J)$ will be the same. The boundaries of these intervals are the memory values at which levels switch from one set to the next. Fig. 16 shows two examples of how the $(H, I, J)$ partition might evolve with two levels.

The algorithm operates in three main steps. In the first step, the sequence of $(H, I, J)$ partitions is determined. For instance, we determine which of the two cases illustrated in Fig. 16 holds. In the second step, the values of the boundaries between the different intervals are calculated; in Fig. 16, those would be $\{M_1, M_2, M_3, M_4\}$ or $\{M'_1, M'_2, M'_3, M'_4\}$. Finally, the third step consists in determining the $M$-feasible partition $(H, I, J)$ based on steps 1 and 2, and based on the value of any given memory $M$. For example, if we are in the first case in Fig. 16 and we are given a memory value between $M_3$ and $M_4$, then we know that $(H, I, J) = (\emptyset, \{2\}, \{1\})$.

Recall from Definition 1 that the $M$-feasible partition is defined by the relation of $(M - T_J + V_I)/S_I$ to the two values $(1/K)\sqrt{N_i/U_i}$ and $(1/d_i + 1/K)\sqrt{N_i/U_i}$, for every level $i$. As long as the same relations are maintained, the same partition is chosen as the $M$-feasible partition. Thus, the values of the boundary ranges must be of one of two forms:

$$\begin{aligned} m_i^{I,J} &= \frac{1}{K}\sqrt{\frac{N_i}{U_i}} \cdot S_I + T_J - V_I; \\ M_i^{I,J} &= \left(\frac{1}{d_i} + \frac{1}{K}\right)\sqrt{\frac{N_i}{U_i}} \cdot S_I + T_J - V_I, \end{aligned}$$

[10]Technically, there will be exactly one level in the set $I$, but it will still get zero memory so this distinction is irrelevant.

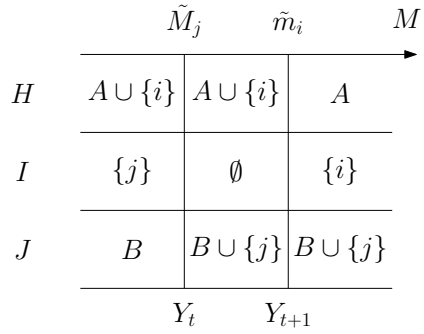


Fig. 17. The only situation where Algorithm 1 results in an interval with $I = \emptyset$. Here, the sets $A$, $B$, $\{i\}$, and $\{j\}$ are disjoint and together form the entire set of levels.

for the apropriate partition $(H, I, J)$. In fact, this partition has to be the one defined by either of the intervals that this boundary divides. For example, $M_2$ in Fig. 16 is the boundary where level 2 moves from $H$ to $I$. Thus, it is of the form $m_2^{I,J}$, where $(H, I, J) = (\emptyset, \{1, 2\}, \emptyset)$, or $(H, I, J) = (\{2\}, \{1\}, \emptyset)$. As it turns out, both choices give the same result.

Let us suppose we are looking at a particular interval with some $(H, I, J)$ partition, and we wish to know its upper boundary. We know that it has to be either $m_h^{I,J}$, $h \in H$, or $M_i^{I,J}$, $i \in I$. Specifically, it has to be the smallest of these values. Since $I$ and $J$ are fixed, this is equivalent to determining the smaller of:

$$\tilde{m}_h = \frac{1}{K}\sqrt{\frac{N_h}{U_h}};$$
$$\tilde{M}_i = \left(\frac{1}{d_i} + \frac{1}{K}\right)\sqrt{\frac{N_i}{U_i}}.$$

But these values do not depend on the chosen partition $(H, I, J)$! Thus, the sequence of $(H, I, J)$ can be uniquely determined solely by these $\tilde{m}_i$ and $\tilde{M}_i$ values.

For what follows, we refer the reader to Algorithm 1 for the various symbols ($x_t$, $Y_t$). It is worth mentioning the following two points. First, the ordering between $x_t$ and $Y_t$ is preserved. That is, $x_t < x_{t'} \implies Y_t < Y_{t'}$. This is ensured by the observation made above that the ordering does not depend on the partition $(H, I, J)$. Second, as mentioned earlier, when applying the function $x \mapsto x \cdot S_I + T_J - V_I$ on some $x = \tilde{m}_i$ or $x = \tilde{M}_i$, it does not matter whether we take the $(H, I, J)$ partition whose interval $x$ lower-bounds or upper-bounds.

Any interval in which $I$ is set to be empty automatically becomes an empty interval. To prove this, we will show that $I = \emptyset$ implies that its upper and lower boundaries match. This is done as follows. Consider the aforementioned interval, and notice that it can only occur if some level moved from $I$ to $J$ at the lower boundary, and another level moved from $H$ to $I$ at the upper boundary. See the bottom case of Fig. 16 for an example, interval $(M_2', M_3')$. Thus, the sequence of $(H, I, J)$ partitions that occurs is of the form shown in Fig. 17.

Let us compute the interval boundaries $Y_t$ and $Y_{t+1}$. Because of our previous observations, we can use $(H, I, J) = (A \cup \{i\}, \emptyset, B \cup \{j\})$ for both. Then,

$$Y_t = \tilde{M}_j \cdot S_I + T_J - V_I = T_J = T_B + N_j/d_j.$$

Similarly,

$$Y_{t+1} = \tilde{m}_i \cdot S_I + T_J - V_I = T_J = T_B + N_j/d_j.$$

Thus, $Y_t = Y_{t+1}$ and the interval is empty.

This concludes the proof of the lemma. ∎

## APPENDIX B
## PROOF OF APPROXIMATE OPTIMALITY FOR THE SINGLE-USER SETUP (THEOREM 6)

As discussed in Section VI-B, we prove Theorem 6 by lower-bounding the optimal rate with expressions of the form shown in (14). Two cases must be considered.

*1) Case $M < 1/6$:* When $M$ is this small, we choose $b = 1$ broadcast message. Recall that, because of regularity condition (2), we have $N_i/K_i \geq 1 > 1/6 > M$. The achievable rate in this case can be upper-bounded by the expression in (15).

Consider now any level $i$. Let $s_i = K_i$. Then,

$$\begin{aligned} v_i &= s_i \left( \min\left\{1, \frac{N_i}{s_i b}\right\} - \frac{M}{b} \right) \\ &= K_i \left( \min\left\{1, \frac{N_i}{K_i}\right\} - M \right) \\ &\geq (5/6) K_i. \end{aligned} \tag{32}$$

We can combine (15) with (32) and (14) to get:

$$\frac{R(M)}{R^*(M)} \leq \frac{6}{5}. \tag{33}$$

*2) Case $M \geq 1/6$:* We will now choose $b = \lceil 6M \rceil \geq 1$.

*a) Bound for $g \in G$:* Consider $s_g = 1$. Then,

$$\begin{aligned} v_g &= \min\left\{1, \frac{N_i}{\lceil 6M \rceil}\right\} - \frac{M}{\lceil 6M \rceil} \\ &\geq \frac{1}{2} - \frac{M}{6M} \\ &= \frac{1}{3}, \end{aligned} \tag{34}$$

because $s_g b = \lceil 6M \rceil \leq 2 \cdot 6M \leq 2N_g$.

*b) Bound for $h \in H$:* Consider $s_h = \lceil K_h/6 \rceil \geq 1$. Then,

$$v_h \geq \frac{K_h}{6} \left( \min\left\{1, \frac{N_h}{4K_h M}\right\} - \frac{M}{6M} \right) = \frac{1}{72} \cdot K_h, \tag{35}$$

because $s_h b = \lceil K_h/6 \rceil \cdot \lceil 6M \rceil \leq 4K_h M \leq 4N_h$.

*c) Bound for $i \in I$:* Consider $s_i = \lceil N_i/6M \rceil$. Then,

$$v_i \geq \frac{N_i}{6M} \left( \min\left\{1, \frac{N_i}{4N_i}\right\} - \frac{M}{6M} \right) = \frac{1}{72} \cdot \frac{N_i}{M}, \tag{36}$$

because $s_i b = \lceil N_i/6M \rceil \cdot \lceil 6M \rceil \leq 4N_i$.

*d) Bound for $J$:* First, if $M \geq N_J$, then the set $J$ contributes nothing to the upper bound on the rate in (12); see (13). Thus we can ignore it, i.e., choose $s_J = 0$ and thus $n_J = 0$ and $v_J = 0$.

So the interesting case is $M < N_J$. Here, we must decode files from multiple levels collectively. Consider $s_J = \lceil N_J/6M \rceil$, where $N_J = \sum_{j \in J} N_j$. Notice that there are enough users and broadcasts to decode all files, because:

$$s_J b \geq \frac{N_J}{6M} \cdot 6M = N_J.$$

However, we must take care that no broadcast considers more than $K_j$ users at a time for any $j \in J$. This can be ensured: since there are $b = \lceil 6M \rceil$ broadcasts, and $b \geq N_j$ for all $j \in J$, then every broadcast need only consider at most one user per level. Hence, all of the $N_J$ files can be decoded, and $n_J = N_J$.

If $M < N_J/6$, we have:

$$\begin{aligned} v_J &= \frac{N_J - s_J M}{b} \\ &\geq \frac{1}{12M} \left( N_J - \frac{N_J}{12M} \cdot M \right) \\ &\geq \frac{144}{11} \cdot \frac{N_J}{M}. \end{aligned} \tag{37}$$

If $N_J/6 \leq M < N_J$, then $s_J$ is actually equal to 1, and:

$$\begin{aligned} v_J &= \frac{N_J - M}{b} \\ &\geq \frac{N_J - M}{12M} \\ &\geq \frac{1}{12} \left( 1 - \frac{M}{N_J} \right). \end{aligned} \tag{38}$$

*3) Multiplicative gap:* By combining (34), (35), (36), (37) and (38) with (12) and (13), and also taking into account (33), we get:

$$\frac{R(M)}{R^*(M)} \leq 72,$$

which concludes the proof of Theorem 6. □

## APPENDIX C
## COMPLETE CHARACTERIZATION FOR THE SMALL EXAMPLE (THEOREM 2)

This appendix proves Theorem 2 in two parts. The first part presents the achievability scheme, and the second part gives the information-theoretic outer bounds.

The rate-memory curve shown in Figure 8—which we are trying to prove is both achievable and optimal—can be described using the following equation:

$$R = \max\left\{3-2M\,,\,\frac{5}{2}-M\,,\,2-\frac{1}{2}M\,,\,1-\frac{M-2}{N_2/2}\right\}. \tag{39}$$

We would like to remind the reader that, in the example considered, the number of less popular files is $N_2 \geq 4$.

### A. Achievability scheme

To prove the achievability of the piece-wise linear curve in Figure 8 and in (39), we need only show the achievability of the corner points, as the rest can be achieved using a memory-sharing scheme between every pair of points. These $(M, R)$ points are:

$$(0,3)\,;\quad \left(\frac{1}{2},2\right);\quad \left(1,\frac{2}{3}\right);\quad (2,1)\,;\quad \text{and}\ \left(2+\frac{N_2}{2},0\right).$$

*1) Point $(M,R)=(0,3)$:* When $M=0$, we do not place any content in the caches, and instead broadcast all the requested files to the users. Since these can be three different files, the peak rate is $R=3$.

*2) Point $(M,R)=(2+\frac{N_2}{2},0)$:* On the other extreme, when $M=2+\frac{N_2}{2}$, the caches are large enough to each store all the popular files, and complementary halves of the unpopular files. Thus, a user can recover any popular file by solely accessing any cache, and can recover any unpopular file by accessing both caches. Therefore, no broadcast is needed, and thus $R=0$ is achievable.

*3) Point $(M,R)=(2,1)$:* In this case, each cache can hold up to two files. Our strategy is to dedicate this memory to the two popular files. Thus, each cache contains all the popular files, but no information about the unpopular files. Therefore, users 1 and 2 can recover their respective requested files without relying on any broadcast, whereas user 3 requires a full broadcast of his requested file. As a result, the rate $R=1$ is achievable.

*4) Point $(M,R)=(1,\frac{3}{2})$:* Let us split each of the popular files into two parts of equal size: $W_1^1=(W_{1a}^1,W_{1b}^1)$ and $W_2^1=(W_{2a}^1,W_{2b}^1)$. We place these in the caches as follows. The first cache contains the first half of each file: $Z_1=(W_{1a}^1,W_{2a}^1)$, whereas the second cache contains the second half of each file: $Z_2=(W_{1b}^1,W_{2b}^1)$. When the user requests are revealed, the BS sends a common message with two parts: $X^{\mathbf{r}}=(X_1^{\mathbf{r}},X_2^{\mathbf{r}})$. The first part directly serves user 3 by giving him the full file he requested: $X_1^{\mathbf{r}}=W_{r_3}^2$. The second part will be $X_2^{\mathbf{r}}=W_{r_1b}^1\oplus W_{r_2a}^1$, which, along with the AP cache content accessed by each of users 1 and 2, allows them to recover their respective file. The BS had to transmit one full file, plus a linear combination of two half files, and, as a result, the total broadcast rate is $R=\frac{3}{2}$.

*5) Point $(M,R)=(\frac{1}{2},2)$:* This case is slightly different in that it requires storing coded content in the caches. We again split the popular files into two halves as before. However, we this time store the following: $Z_1=W_{1a}^1\oplus W_{2a}^1$; and $Z_2=W_{1b}^1\oplus W_{2b}^1$. The BS will transmit the following broadcast: $X^{\mathbf{r}}=\left(W_{r_3}^2,W_{r_1b}^1,W_{r_2a}^1\right)$. This will serve all the requests of the users. The broadcasts message consists of a full file and two half-files, and thus the total rate is $R=2$.

### B. Outer bounds

We will now prove that the above scheme is optimal with respect to information-theoretic bounds. We do that by showing that the rate is larger than each one of the expressions in the maximization in (39).

In all of the following, inequalities marked by $(*)$ are due to Fano's inequality, and the $\varepsilon_F$ term that arises along with these inequalities is a term that decays to zero as $F\to\infty$.

*1) First expression:* Let the request vector be $\mathbf{r} = (1, 2, 1)$, and consider both caches along with the broadcast $X^{\mathbf{r}}$. Then, for all $F$,

$$\begin{aligned} RF + 2MF &\geq H\left(Z_1, Z_2, X^{\mathbf{r}}\right) \\ &= H\left(Z_1, Z_2, X^{(1,2,1)} \middle| W_1^1, W_2^1, W_1^2\right) \\ &\quad + I\left(W_1^1, W_2^1, W_1^2; Z_1, Z_2, X^{(1,2,1)}\right) \\ &\overset{(*)}{\geq} 3F(1-\varepsilon_F) \\ R + 2M &\geq 3(1-\varepsilon_F). \end{aligned}$$

By taking $F \to \infty$, we get:

$$R \geq 3 - 2M.$$

*2) Second expression:* Consider the two request vectors $\mathbf{r}_1 = (1, 2, 1)$ and $\mathbf{r}_2 = (2, 1, 2)$.

$$\begin{aligned} 2\left(RF + MF\right) &\geq H\left(Z_1, X^{\mathbf{r}_1}\right) + H\left(Z_2, X^{\mathbf{r}_2}\right) \\ &= H\left(Z_1, X^{\mathbf{r}_1} \middle| W_1^1\right) + I\left(W_1^1; Z_1, X^{\mathbf{r}_1}\right) \\ &\quad + H\left(Z_2, X^{\mathbf{r}_2} \middle| W_1^1\right) + I\left(W_1^1; Z_2, X^{\mathbf{r}_2}\right) \\ &\overset{(*)}{\geq} H\left(Z_1, Z_2, X^{\mathbf{r}_1}, X^{\mathbf{r}_2} \middle| W_1^1\right) + 2F(1-\varepsilon_F) \\ &= H\left(Z_1, Z_2, X^{\mathbf{r}_1}, X^{\mathbf{r}_2} \middle| W_1^1, W_2^1, W_1^2, W_2^2\right) \\ &\quad + I\left(W_2^1, W_1^2, W_2^2; Z_1, Z_2, X^{\mathbf{r}_1}, X^{\mathbf{r}_2} \middle| W_1^1\right) \\ &\quad + 2F(1-\varepsilon_F) \\ &\overset{(*)}{\geq} 5F(1-\varepsilon_F) \\ 2R + 2M &\geq 5(1-\varepsilon_F). \end{aligned}$$

By taking $F \to \infty$, we get:

$$R \geq \frac{5}{2} - M.$$

*3) Third expression:* Consider the following four request vectors: $\mathbf{r}_1 = (1, 2, 1)$, $\mathbf{r}_2 = (2, 1, 2)$, $\mathbf{r}_3 = (1, 2, 3)$, and $\mathbf{r}_1 = (2, 1, 4)$. For clarity, define $\mathcal{X} = (X^{\mathbf{r}_1}, X^{\mathbf{r}_2}, X^{\mathbf{r}_3}, X^{\mathbf{r}_4})$ and $\mathcal{W}^2 = (W_1^2, W_2^2, W_3^2, W_4^2)$.

$$\begin{aligned} 2(2RF + MF) &\geq H\left(Z_1, X^{\mathbf{r}_1}, X^{\mathbf{r}_2}\right) + H\left(Z_2, X^{\mathbf{r}_3}, X^{\mathbf{r}_4}\right) \\ &= H\left(Z_1, X^{\mathbf{r}_1}, X^{\mathbf{r}_2} \middle| W_1^1, W_2^1\right) \\ &\quad + I\left(W_1^1, W_2^1; Z_1, X^{\mathbf{r}_1}, X^{\mathbf{r}_2}\right) \\ &\quad + H\left(Z_2, X^{\mathbf{r}_3}, X^{\mathbf{r}_4} \middle| W_1^1, W_2^1\right) \\ &\quad + I\left(W_1^1, W_2^1; Z_2, X^{\mathbf{r}_3}, X^{\mathbf{r}_4}\right) \\ &\overset{(*)}{\geq} H\left(Z_1, Z_2, \mathcal{X} \middle| W_1^1, W_2^1\right) \\ &\quad + 2 \cdot 2F(1-\varepsilon_F) \\ &= H\left(Z_1, Z_2, \mathcal{X} \middle| W_1^1, W_2^1, \mathcal{W}^2\right) \\ &\quad + I\left(\mathcal{W}^2; Z_1, Z_2, \mathcal{X} \middle| W_1^1, W_2^1\right) \\ &\quad + 4F(1-\varepsilon_F) \\ &\overset{(*)}{\geq} 8F(1-\varepsilon_F) \\ 4R + 2M &\geq 8(1-\varepsilon_F). \end{aligned}$$

By taking $F \to \infty$, we get:

$$R \geq 2 - \frac{1}{2}M.$$

*4) Fourth expression:* Consider the $N_2$ request vectors $\mathbf{r}_1, \ldots, \mathbf{r}_{N_2}$, such that:

$$\begin{aligned}\mathbf{r}_k &= (1,2,k) \qquad \forall k \leq \left\lfloor \frac{N_2}{2} \right\rfloor, \\ \mathbf{r}_k &= (2,1,k) \qquad \forall k \geq \left\lfloor \frac{N_2}{2} \right\rfloor + 1.\end{aligned}$$

Define $\mathcal{X}^e = \{X^{\mathbf{r}_k} : \ k \text{ is even}\}$ and $\mathcal{X}^o = \{X^{\mathbf{r}_k} : \ k \text{ is odd}\}$. Thus, each of $\mathcal{X}^e$ and $\mathcal{X}^o$ contains at least one broadcast $X^{\mathbf{r}_k}$ such that $\mathbf{r}_k = (1,2,k)$, and one broadcast $X^{\mathbf{r}_l}$ such that $\mathbf{r}_l = (2,1,l)$. Furthermore, define, for clarity, $\mathcal{W}^2 = (W_1^2, \ldots, W_{N_2}^2)$. Then,

$$\begin{aligned}N_2RF + 2MF &\geq H\left(Z_1, \mathcal{X}^e\right) + H\left(Z_2, \mathcal{X}^o\right) \\ &= H\left(Z_1, \mathcal{X}^e \middle| W_1^1, W_2^1\right) \\ &\quad + I\left(W_1^1, W_2^1; Z_1, \mathcal{X}^e\right) \\ &\quad + H\left(Z_2, \mathcal{X}^o \middle| W_1^1, W_2^1\right) \\ &\quad + I\left(W_1^1, W_2^1; Z_2, \mathcal{X}^o\right) \\ &\overset{(*)}{\geq} H\left(Z_1, Z_2, \mathcal{X}^e, \mathcal{X}^o \middle| W_1^1, W_2^1\right) \\ &\quad + 4F(1-\varepsilon_F) \\ &= H\left(Z_1, Z_2, \mathcal{X}^e, \mathcal{X}^o \middle| W_1^1, W_2^1, \mathcal{W}^2\right) \\ &\quad + I\left(\mathcal{W}^2; Z_1, Z_2, \mathcal{X}^e, \mathcal{X}^o \middle| W_1^1, W_2^1\right) \\ &\quad + 4F(1-\varepsilon_F) \\ &\overset{(*)}{\geq} (4+N_2)F(1-\varepsilon_F) \\ N_2R + 2M &\geq (4+N_2)(1-\varepsilon_F).\end{aligned}$$

By taking $F \to \infty$, we get:

$$R \geq \frac{4+N_2-2M}{N_2} = 1 - \frac{M-2}{N_2/2}.$$